Emergence of Things Felt: Harnessing the Semantic Space of Facebook Feeling Tags

Completed Research Paper

Chris Zimmerman
Computational Social Science Lab
ITM- Copenhagen Business School
Howitzvej 60, Frederiksberg 2000
Denmark
cz.itm@cbs.dk

Mari-Klara Stein
Department of IT Management (ITM)
Copenhagen Business School
Howitzvej 60, Frederiksberg 2000
Denmark
mst.itm@cbs.dk

Daniel Hardt
Department of IT Management (ITM)
Copenhagen Business School
Howitzvej 60, Frederiksberg 2000
Denmark
dh.itm@cbs.dk

Ravi Vatrapu
Computational Social Science Lab
ITM - Copenhagen Business School
Howitzvej 60, Frederiksberg 2000
Denmark
rv.itm@cbs.dk

Abstract

In 2013 Facebook launched a feature allowing users to add a feeling tag to their posts as part of their daily interactions online. Our research leverages the text accompanying all such volunteered feeling tags in an effort to map the semantic space of 'Facebook feelings' as they are catalogued by the crowd. By letting the data speak for itself, a folksonomy of feelings reveal temporal and social patterns in the most commonly shared feelings. Unlike many such studies, however, we do not only focus on examining the patterns emerging from big data, but also put the expressed feelings to work using machine learning towards both the classification and detection of emotions. This paper first demonstrates that feelings expressed online self-organize along the same lines (valence and arousal dimensions) experts in psychology and emotions have organized them for decades. As we enter the debate of classifying human emotions, our analysis directly contrasts Facebook's manifestation of feelings with prior theoretical proposals to detect both similarities and differences from past assumptions. In line with the 'exhibitional' nature of Facebook, we illustrate that 'extreme' feelings, such as excitement and anger, are expressed in even more extreme levels of both valence and arousal. Beyond contrasting the folksonomy of feelings with dimensional mappings of emotions proposed by past research, we further utilize artificial intelligence techniques towards building a test version of an automatic 'Feelings Meter' able to detect feelings from text.

Keywords: Facebook, Feelings, Sentiment Analysis, Arousal, Social Media, Marketing

Introduction

On April 13th, 2013, Facebook launched a new feature allowing users to add a 'feeling' tag to their posts. Users of the social network had previously been allowed to add photos, links, as well as tag their friends and locations to complement their text in a status update. From this point forward users have also had the *option to choose from roughly 120 different pre-defined feelings to add to their status updates at any given time.* The addition of these feeling tags (a form of annotation) has provided a mechanism for users to augment their posted status updates with an expressed feeling. As of January 2014, 71% of online adults use Facebook (Pew Research Center 2015), making it the world's most widely utilized social networking site. By the end of August 2015, the company reported over 1 billion users in a single day – approximately 1 in every 7 people on the planet (BBC News 2015). As such, Facebook can provide a unique insight into a large set of posts *tagged with feelings,* offering unprecedented access to the expressed emotional lives of the public (see also Kamvar and Harris, 2011).

The purpose of this research is, *first*, to *understand* these feelings that users choose to explicitly tag and publicly share. It is well-known that the meaning of many familiar concepts on social media is not always the same as outside of social media. For example, the concepts 'friend' and 'like' have taken on quite a distinct connotation on Facebook (Hogan 2010). Accordingly, the semantic space of these concepts is not the same as outside of social media. Given that one can feel 'sexy', 'pissed', tipsy' or 'awesome' (notably not the same as feeling 'awe') on Facebook, we find it important to map the semantic space of 'Facebook feelings'. In other words, we aim to understand the basic patterns of how feelings are shared on Facebook and how they can be described in terms of valence (positivity, negativity) and levels of arousal. Furthermore, we explore how (if at all) do the user-categorized 'Facebook feelings' differ, on the valence and arousal dimensions, from previously theorized mappings of feelings (Russell, 1983; Scherer 2005). In that regard, our paper aligns with recent works that increasingly argue for the treatment of online and offline phenomena as potentially *different*, but equally *valid* (Ellis and Tucker 2015). As social media become a new forum for the production of emotional activity, it is, in our view, essential to recognize that an online experience (e.g., 'digital emotion') is not inferior to or less valid than an offline experience (ibid.). Using face-to-face communication and 'real' emotions as the yardstick, compared to which all online experiences are seen as less rich, does not account for the very real and rich *extensions* to our experiential worlds that social media and digitalization have brought about. Our paper is not "just about defining [a technological space] that people can experience emotions *within*", but rather about how Facebook is "allowing people to produce new and innovative emotional solutions" (Ellis and Tucker, 2015, p. 178). In sum, we first set out to gain a better understanding of the basic patterns of how feelings are expressed on Facebook. *Second,* our aim is to *inform organizational practices* related to social media analytics (Holsapple et al. 2014), particularly *sentiment* analysis (cf. Stieglitz and Dang-Xuan 2013). A better understanding of user-categorized feelings also allows us to build better analytics tools that are able to process data on a more granular level and reveal more about user sentiment than just its polarity in terms of positivity or negativity. We present and discuss a test version of such a tool in the practical implications section.

The dataset we collected from Facebook is comprehensive in its size and scope, capturing almost 12 million Facebook posts and covering all public instances that included a feeling tag since the feature was introduced (over 18 months). The data itself is of particular interest in that tags are deliberate user annotations of feelings relating to each status update. It also includes the additional contextual tags provided by users, lending more insight into which feelings are co-shared with other (tagged) people, and when feelings are cross-tagged at locations. Unlike Facebook's own emotional contagion study (Kramer, et al. 2014), our data collection of public posts had neither privileged access to Facebook's user profiles and private posts nor did we conduct digital field experiments/interventions with users' Facebook feeds. Instead, our analysis used natural language processing techniques leveraging classifiers that focus directly on the language used when status updates declare a corresponding emotional state.

Our findings reveal the most commonly shared Facebook feelings and their temporal and social patterns. Furthermore, the classifiers allow the dataset to self-organize, revealing the most prominent feelings as well as their positions in the two-dimensional valence-arousal space. Being able to compute the valence and arousal positioning of user-categorized 'Facebook feelings' also provides a comparison point to previous research that has argued over the conceptual categorization of feelings (cf. Russell 1983; Scherer 2005). For example, our analyses indicate the possibility that 'extreme' feelings, such as *excitement*,

anger and *sadness*, tend to be expressed in an even more extreme manner in terms of both valence and arousal on Facebook. Furthermore, we find that Facebook provides users with both an additional socio-technical space in which to express feelings (Ellis and Tucker 2015), and also a space where novel emotional scripts (Ashkanasy 2003) are created and honed. Our contributions in this paper are threefold: domain-specific, methodological and practical. First, our paper expands the understanding of user-categorized feelings on social media. To our knowledge our study is one of the first to (a) study feelings explicitly expressed through tags on Facebook, and (b) map feelings expressed on social media both in terms of valence and arousal. Prior research on feelings expressed on social media has largely focused on Twitter and blogs (Kamvar and Harris 2011; Barnaghi et al. 2015) and has utilized traditional sentiment analysis that focuses only on valence (positive-negative polarity or subjectivity ratios) (Pang et al. 2002; Asur and Huberman 2010). Methodologically, our study draws on a unique data set, taking advantage of insights that can be generated by 'big data' and applies natural language processing (NLP) techniques to the study of discrete feelings. While 'big data' studies and NLP techniques are popular in traditional sentiment analysis (Barnaghi et al. 2015), most studies of discrete feelings have to date relied on 'small data' and experimental or qualitative methods (cf. Scherer, 2005). Thus, our study is able to provide important evidence with regard to patterns in feeling expression emerging from the crowd on social media. Lastly, based on the understanding gained we are also able to propose a test version of a practical analytics tool and briefly outline scenarios for its potential application in the area of corporate branding.

The remainder of the paper is organized as follows. Next we define the concepts of emotions and feelings and present relevant theories from psychology and related work on social media and emotions in information systems. The data set section provides details on the collection, processing and analysis of the data. Computational linguistics aspects of the classifiers are discussed next in the method section. Finally, results are presented in the findings section with the discussion of their substantial interpretation, relations to extant literature and practical implications considered last.

Theoretical Background

In order to better understand the semantic space of 'Facebook feelings', we draw on extant research, first, to define what we mean by the terms emotion and feeling and, second, to outline the dimensional approach for describing and measuring feelings (Russell 2003; Scherer 2005). We also review prior research that has explored emotions and feelings in social media contexts in particular.

Definitions: Feelings and Emotions

In recent years, emotion has become an increasingly popular topic in organization studies (e.g. Benozzo and Colley 2012; Grant 2013; Lindebaum and Cartwright 2010) as well as in information systems (IS) research (Bagozzi 2007; Beaudry and Pinsonneault, 2010; McGrath 2006; Ortiz de Guinea and Markus, 2009; Stein et al. 2014; Stein et al. 2015; Thompson 2012; Zhang 2013). Despite this rise in interest, emotions still constitute a very demanding research object (Kopelman et al. 2006). First, different theoretical traditions discuss a multitude of emotion-related concepts, such as feelings, moods, affect, and temperament and often disagree on their definitions (Barsade and Gibson 2007; Scherer 2005). Furthermore, many data collection situations may only reveal the display of emotions, not the internal experience of emotions, thus, partially obscuring the phenomenon under study (Kopelman, et al. 2006).

In this paper, we define emotion as an "*episode of interrelated, synchronized changes in the states of all or most of the five organismic subsystems (cognitive, neurophysiological, motivational, motor expression and subjective feeling) in response to the evaluation of an external or internal stimulus event as relevant to major concerns of the organism*" (Scherer 2005: 697). In other words, this means emotions "arise as reactions to situational events in an individual's environment that are appraised to be relevant to his/her needs, goals, or concerns. Once activated, emotions generate subjective feelings … motivational states with action tendencies, arouse the body with energy-mobilizing responses … and express the quality and intensity of emotionality outwardly and socially to others" (Zhang 2013: 251). This definition allows for an important distinction between (subjective) feelings and emotions. According to Scherer (2005), subjective feeling captures only one component of emotions - the subjective experience of it. Most existing IS and organizational research (Beaudry and Pinsonneault 2010; Russell 1983; Stein et al. 2015) does not make this distinction. In this study, we specifically look only at subjective feelings, and use the term 'feeling(s)' and 'subjective feeling(s)' interchangeably. Given that a feeling by definition is a subjective experience, one option to understand feelings is to ask individuals to report on the nature of

their experience. This aligns well with our context of studying user-categorized feelings on Facebook. It is beyond the scope of this paper to go into the discussion of possible incongruences between expressed feelings and actually felt feelings, and issues of display/feeling rules and emotion work (Fineman 2008).

Measuring Feelings - Dimensional Approach

Feelings can be described and measured in terms of a number of different dimensions, such as intensity, duration, valence, arousal and tension (Scherer, 2005). A dimensional approach to describing subjective feelings was pioneered by one of the founders of modern psychology, Wilhelm Wundt in 1905 (ibid.). He suggested that, through introspection, individuals are able to describe their feelings by positioning them in a three-dimensional space of valence (positive–negative), arousal (calm–excited), and tension (tense–relaxed) (Scherer 2005). The two-dimensional adaptation of this idea, retaining only the valence and arousal dimensions[1] has become widely accepted in emotions research (Russell 1983; Russell 2003; Barsade and Gibson 2007). The dimension of tension is often excluded due to difficulties in consistently identifying what the dimension describes: tension, control, or potency (ibid.).

Consequently, a feeling can be described as an experience "that is an integral blend of hedonic (pleasure–displeasure) and arousal (sleepy–activated) values" (Russell 2003, p. 147). Any feeling, thus, can be described as a point in the valence-arousal space (ibid.; Scherer 2005). The valence dimension (pleasure–displeasure or positive-negative feeling), ranges from the negative extreme (e.g., sad or depressed) through a neutral point to the positive extreme (e.g., happy). The arousal dimension ranges from the low arousal end (e.g., sleepy, calm or tired) to high arousal end (e.g., excited or angry), through various stages of alertness. One such two-dimensional map is provided in Figure 1.
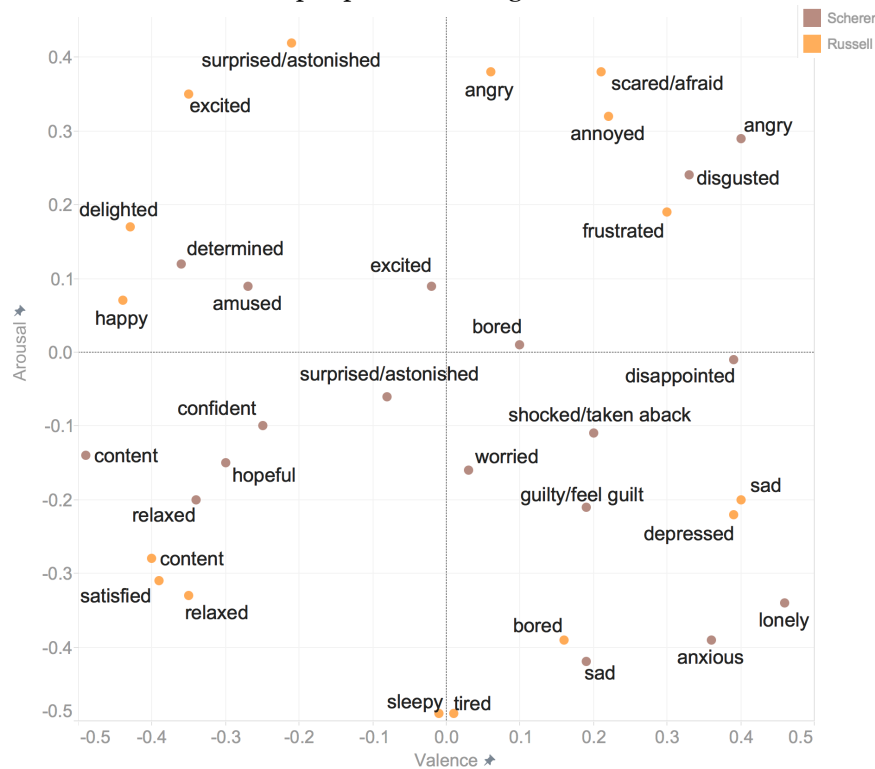

Figure 1. Commonly Accepted Two-Dimensional (Valence, Arousal) Semantic Space for Feelings (adapted from Russell 1983; Scherer 2005)

Figure 1 combines findings from two well-known theoretical sources - Russell (1983) and Scherer (2005). Russell (1983) demonstrated that a number of feeling-related terms (whether judged by Gujarati,

---

[1] While the valence (positive to negative) and arousal (aroused to calm) dimensions are the most commonly used, it is important to note that they do not cover all variation in the experience of feelings. For example, intense anger may be a high arousal feeling, while intense sadness may be a low arousal feeling (Scherer, 2005). The dimension of intensity, thus, does not correspond to arousal.

Croatian, Japanese, Chinese or English speakers) fell in a roughly similar circular order in the two-dimensional space. Scherer (2005) compared Russell's results to his own study of German terms, and found broad similarities so that most of the equivalent feeling terms fell in the same quadrants[2]. Having now introduced the key concept of feeling, as well as discussed how feelings can be described in a two-dimensional semantic space, we next consider feelings in the context of social media specifically.

Feelings and Social Media

Information and communication technologies (ICTs) are implicated in most emotion-related processes – ICTs can stimulate emotional responses, while the adoption and use of ICTs is, in turn, impacted by emotions and feelings (Stein et al. 2015; Zhang 2013; Beaudry and Pinsonneault 2010). Social media in particular, such as Facebook, have been shown to not only facilitate but also influence the generation of feelings, for example, through emotional contagion processes (Kramer et al. 2014).

Many businesses and researchers alike, do not underestimate the value of understanding the troves of data generated by users on Facebook and other social media (Culnan et al. 2010; Mandviwalla and Watson 2014; Vatrapu 2013). The field of social media analytics (Holsapple et al. 2014) has emerged as a result, involving the development and evaluation of "informatics tools and frameworks to collect, monitor, analyze, summarize, and visualize social media data" (Abrahams et al. 2013: 872).

Our unique dataset allows for a wide range of such social media analytics efforts (including sentiment, affect, and semantic analyses, cf. Abbasi and Chen 2008; Chau and Xu 2012; Mukkamala et al. 2015; Stieglitz and Dang-Xuan 2013; Zimmerman et al. 2014). In this paper, however, we take a first step towards understanding the phenomenon of Facebook feeling tags itself. Given the relative novelty of the feature that allows people to express and communicate their feelings on Facebook, we set out to first map the semantic space of 'Facebook feelings'. We contend that in order to build useful analytics tools and conduct effective analyses, it is essential to understand the data in depth from relevant domain-specific perspectives. This is particularly the case in the world of social media, where the semantics of well-known concepts is not always the same as outside of social media (Hogan 2010).

Data Set

The data set used in this paper captures feeling-tagged public posts on Facebook. Despite access limitations, there are several advantages to data collected from Facebook.

Platform Selection – Advantages and Limitations

The most fundamental trade-off of collecting data from Facebook is one of adoption versus ease of collection. Facebook has become one of the most ubiquitously inhabited social media platforms and is a part of people's daily lives. While all social media channels have a sample bias problem (race, gender, geographic adoption), Facebook has over 1.44 billion monthly active users (MAU) and is widely used in most of the world ('Facebook Company Info - Statistics,' 2015). Twitter, conversely, is only prominently used in a handful of countries, with much lower adoption rates overall (Pew Social Media Report 2015). Yet research to date that examines emotional texts has largely been conducted on blogs or micro-blogging platforms (Golder 2011; Kamvar and Harris 2011; Marsella and Gratch 2014). For example, Golder (ibid.) identified patterns for daily mood cycles by applying traditional sentiment analysis to over 500 million tweets. In fact, much of the past research into online emotions has relied on data from Twitter instead of Facebook for several reasons. The first is ease of access. Over 90% of Twitter users have their profile set to public, whereas less than 50% do on Facebook (Tufekci 2014). Tweets are short in nature, publicly visible within a directed but open network graph. Tufekci (2014) claims that such advantages have led to Twitter becoming the "Drosophila melanogaster" or the model organism for social media papers, while the trade-offs of this are rarely discussed. Yet quantity and accessibility of data do not necessarily guarantee quality of data for the purposes of specific emotional mapping. For example, when measuring sentiments on Twitter, one could question how many expressed feelings are actually emanating from text created or re-tweeted by 'bots' (Ferrara, et al., 2015).

---

[2] Scherer's (2005) study included a wider range of feelings than Russell's (1983); we have included some of the feelings only considered in Scherer (ibid.) in Figure 1 to allow for a later comparison with our Facebook data. Figure 1 represents our digitized version of a similar figure provided in Scherer (ibid. p. 720).

Personal Nature. Despite the technical advantages for data collection from open online social networks like Twitter, analytical advantage with Facebook data lies in the intentions of the user who is sharing information. When content is shared on LinkedIn the perceived audience is often a professional network. Twitter, for many users, is a public forum where anyone who wants to can listen, thus, mixing private and professional connections. Because Facebook requires mutual connections of a personal nature this audience is mostly people we know and care about. The intended receiver of posts is friends and family, thus, making Facebook one of the most suitable online domains for investigating personal feelings.

Contextual Disambiguity. Given our specific methodological goal of social text analysis from big data, another advantage of the Facebook platform is the existence of tags in status updates. Thus, when Facebook users add a feeling tag to their post, this constitutes a definitive emotional categorization, despite the potential existence of irony or sarcasm within the text. The advantage of tagging on Facebook also lies within the contextually unambiguous usage of this feature. For example, when hashtags or mentions of feelings are used discursively in the text of other social media posts, they may be describing another person, or a feeling from another point in time. The surrounding text may completely change the overall intended meaning of the post as a whole. Figures 2a and 2b below illustrate the difference in contextual clarity between a *discursive* (2a) and a *tagged* post (2b), both mentioning feeling happy.[3] The former contains a certain degree of negativity in the text, but also the words 'feeling happy' in the body of the text. Whereas the deliberately tagged post has 'feeling happy' annotating the entire post. The post text then contains terms that would benefit in training a classifier in recognizing words and characters (e.g., exclamation marks) that frequently correspond to declarations of 'feeling happy' (the same applies to other feelings of course). In summary, the selection of the Facebook platform for data collection was because of its global adoption, personal nature, and unique volunteer mechanisms that allow for the best opportunity at capturing the language associated with user-categorized feelings online.

Figures 2a and 2b. Example of Discursive vs. Tagged Posting.

Data Collection

Preliminary observations of Facebook feelings listed over a month-long period provided a list of 143 possible feelings, which were then used as search strings. The order of emotions presented to Facebook users fluctuates, and only 120 are visible at any given time.[4] This list provided search strings for data collection aimed at capturing all English-language instances of 'Feeling (X, Y, Z)'.

In October 2014, the Radian6 tool was used to download all public mentions of 'Feeling X, Y, Z' based on the 143 different combinations observed initially[5]. A full 18 months of data was thus collected to capture the entire history of activity since the introduction of the feelings tag feature in status updates. The resulting dataset was then visualized and analyzed using Tableau Desktop[6]. Tableau was used to first explore and understand the dataset, and later to examine results from text classifiers.

Several limitations exist when collecting social big data for research purposes. A fundamental limitation is that of users' default and post-specific privacy settings. This study has had no privileged or direct access to data from Facebook itself. Thus, we rely on posts where the post privacy setting is set to 'public'. One may speculate in the differing behavior between those who post publicly and those who post exclusively within

---

[3] Note: both posts are in the public domain of Facebook accessibility, as distinguished by the globe icon next to the date.
[4] Systematic weekly captures over six months have shown that there are at least 155 different pre-defined options presented to users, yet about 120 are visible at any given time. Locations specific to the user, and seasonal popularity of feelings by all Facebook users may impact the order, however, the feature-specific details are only known to Facebook.
[5] Radian6 is an enterprise software tool for social media analytics and response management, and is part of the Salesforce marketing cloud solution suite. Historical data was collected with special permission from Salesforce. http://www.radian6.com
[6] Tableau is a visual analytics software tool for data visualization and business intelligence. All visual data representations in this paper were generated in Tableau v9. http://www.tableau.com

their 'walled garden'.[7] However the sheer volume of persons who do voluntarily (or involuntarily by default setting) share their posts with feeling tags have allowed us to collect over 1.6 million posts.

Data Preparation

Data Filtering. All public posts using any of the 143 observed feelings were downloaded ('feeling X,Y,Z'). The post-level data included plain text status updates, photos, videos, and shared links, yielding a total of *11,908,715 total results posted by 8,177,586 unique actors*. These include those posts made with feeling tags and those without (discursive mentions of the same feelings), given that the Radian6 tool does not offer a way to pre-filter for results with the tag feature.[8] Certain textual trace elements indicated whether a post contained a feeling tag, or simply a discursive mention. These included the punctuation dividing text and tag as well as the proximity to the end of the post. Such patterns were initially identified within the raw data. Corresponding regular expressions were then used to effectively filter and mine the dataset for examination of the feeling-tagged posts. Of the nearly 12M posts downloaded, 86.4% of posts included discursive mentions of a feeling, leaving *1,618,499 feeling-tagged posts* for usage in our computational linguistics analysis. Figure 3 outlines stages in the data handling process, as well as the upcoming applications of our research.
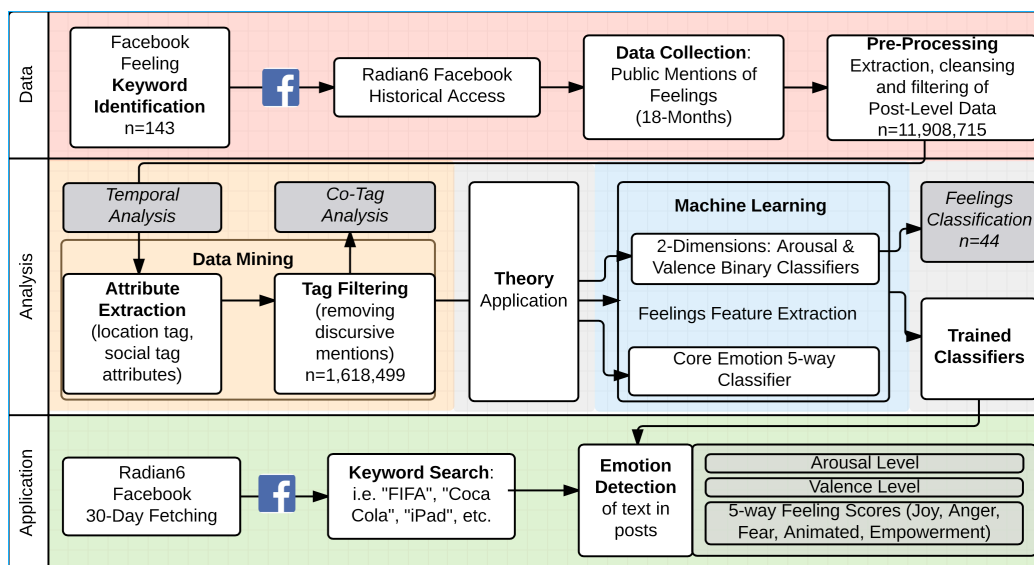

Figure 3. Process diagram depicting the flow of this study from data assemblage to analysis to re-application of a trained model on new testing data.

Data Provenance. When acquiring social data of this size, there is a high level of noise (as well as spam) within millions of online postings. The first issue that we addressed was that of duplicates. Second, a noticeable amount of non-English language data existed from users who have their Facebook profile set up in English, and thus use English tag labels, even though they are posting in another language. Interestingly, mixed use of language was often the case in these instances. These latter instances were not filtered (given the size in comparison to that of remaining English posts) and all words were treated equally in their ability to inform language classifiers. In our observations of thousands of posts, we found that much of the noise in the dataset was distilled out when filtering out discursive mentions of feelings and focusing on the tagged feelings subset.

Delimitations and Selection for Analysis. As noted earlier, more than 143 different Facebook feelings are shared by users. The long tail observed in our dataset suggests that many of these 143 tags would likely yield very low volumes from the limited scope of public posts. Thus, our analysis proceeded to drill down to the most relevant subset of feelings for deeper analysis. We manually selected 44 feelings

---

[7] This is a valid criticism were this research to focus on the social graph (demographics and usage patterns) of the data like most other studies do, rather than simply employing the social text itself.
[8] It is unclear how the Facebook company chose these 143 emotions to be offered in its 'feelings' list and may be extracted from the usage patterns by users.

for comparison with existing dimensional classifications based on three criteria. The first and primary criterion was significant volume of posts per each feeling tag. The selected feelings were among those with the highest volume of source data (averaging 19.7 thousand per feeling). The second criterion was congruence with existing typologies. The majority of feelings (26/44) overlapped with those from existing typologies, such as those of Scherer (2005) and Russell (1983). These two criteria identified the same feelings on some occasions (e.g., feeling 'happy'), and different on others (e.g., feeling 'awesome' has significant volume on Facebook, but is not present as such in extant emotion typologies). Third, our selection attempted to balance relatively equal distribution across the dimensions of valence and arousal.

Method: Natural Language Processing

Standard approaches to sentiment analysis (Pang and Lee 2000) involve the following steps, which are common to many text classification tasks: (1) Preprocessing: text is tokenized (so that words can be separately identified), and often lower-cased; (2) Feature Extraction: identifying single words (unigrams) as well as two (bigrams) and three word (trigrams) sequences; (3) Classification: a supervised machine learning algorithm is selected, which is able to determine which combinations of features best predict the classification of interest. We used the Maximum Entropy algorithm, which has shown good performance on sentiment analysis and related applications (Pang et al. 2002). We followed the above process in all the classification tasks described in this paper.

Valence and Arousal Classifiers. A standard approach to measuring feelings in psychology involves two dimensions: *valence* and *arousal*. We built a classifier for each of these two dimensions, in the following way. First, we selected feelings that indicated extreme degrees of arousal or valence, as shown in Tables 1 and 2. For example, sad, disgusted and disappointed are theoretically considered to be the most negative feelings, while happy and wonderful are considered the most positive (Scherer 2005). Similar principle was applied when selecting the feelings to indicate low arousal and high arousal. Based on prior research (Russell 1983; Scherer 2005), we thus presume that the feelings given in Tables 1 and 2 represent the extremes and will place accordingly on the two-dimensional valence-arousal map for our Facebook feelings dataset. The placement of other feelings will be determined purely from the data. Furthermore, while we presume, for example, that sleepy will place in the low end of the arousal dimension, the data will reveal how it places on the valence dimension. Additionally, we considered the volume of available data (number of Facebook postings tagged with that particular feeling), as the accuracy of a classifier depends on the size of the training data set.

Table 1. Negative and Positive Valence Feelings			
Negative Valence	**Number of Postings**	**Positive Valence**	**Number of Postings**
Sad	70,835	Happy	114,258
Disgusted	1,565	Great	55,179
Disappointed	2,533	Wonderful	54,690
Total	74,933	**Total**	224,127

Table 2. Low and High Arousal Feelings			
Low Arousal	**Number of Postings**	**High Arousal**	**Number of Postings**
Tired	13,431	Excited	155,290
Relaxed	3,747	Angry	12,679
Sleepy	4,187	Pissed	3,850
Total	21,365	**Total**	171,819

Both the arousal and valence data sets are extremely imbalanced, reflecting what may be a general tendency of Facebook users towards posting status updates that are tagged with high arousal, positive feelings (e.g., excitement). To address this issue, we constructed balanced versions of each data set, by randomly selecting 21,365 high arousal postings (to match the total volume of low arousal postings), and 74,933 positive valence postings (to match the total volume of negative valence postings).

With these balanced data sets, we used Maximum Entropy to produce binary classifiers, using 10-fold validation. Maximum Entropy is a classification algorithm frequently used for sentiment analysis and

other text classification tasks. It learns values for feature weight parameters – in our setting, the features are unigrams, bigrams and trigrams. These are standard features for sentiment analysis as well as other applications in natural language processing (NLP). The systems were trained and tested using Mallet, a Machine Learning Toolkit for NLP (http://mallet.cs.umass.edu). The training was performed using 10-fold cross validation. This is a standard technique in which the data is partitioned into a test set consisting of 1/10 of the data and a training set consisting of the remaining data. This is done 10 times for 10 different randomly determined partitions. Test and training results are given below (Table 3)[9].

Table 3. Classifier Test and Training Accuracies		
Classifier	**Training Accuracy**	**Test Accuracy**
Arousal	.9969	.8013
Valence	.9964	.8338

In addition to test and training accuracies, the binary ordering of feelings can also be interpreted manually. The results from running the binary classifiers on the data set of the selected 44 'Facebook feelings' are shown in Figure 4. All feelings are in rank order for both arousal and valence. Each of these can be interpreted by how much arousal/passivity and positivity/negativity they intuitively include by definition in comparison to other feelings in the spectrum. For example, feeling 'disappointed' is more negative than 'angry', while 'angry' shows a higher level of arousal. Given the demonstrated high degrees of accuracy, we chose to use the two classifiers in our two-dimensional assessment of Facebook feelings.
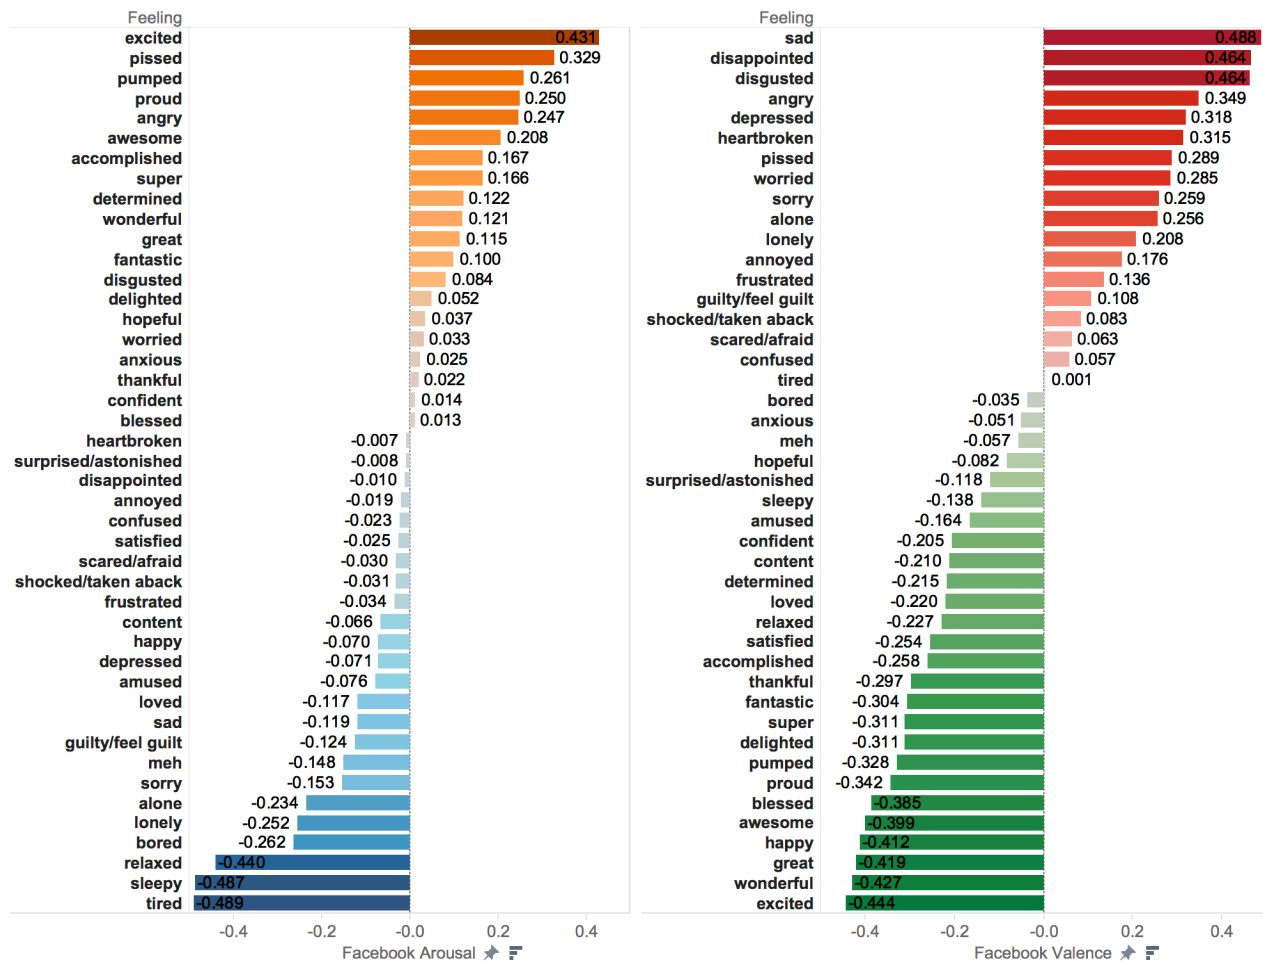

Figure 4. Results: Binary Arousal and Valence Classifiers on 44 Facebook feelings

9 Training accuracy demonstrates the accuracy of the classifier when tested on the same data it was trained on. This number is inherently always higher than the 'real' classifier accuracy. The test accuracy results should, thus, give a better sense of how accurate the classifier is.

Findings

In the following sections we present the results from our analyses. The 18 months worth of posts yield insights into basic tendencies of which feeling tags are used and when, as well as reveal contextual patterns for where certain feelings are expressed and with whom. Second, after explaining user activity patterns, we consider the results of the trained classifiers to draw the two-dimensional (valence-arousal) space for 'Facebook feelings'.

Basic Patterns – What and When?

In general, the dataset revealed that positive feelings are more commonly shared on Facebook than negative feelings. Individually, feeling 'excited' was the most publicly shared emotion (155,290 tags out of 582,691 total mentions). Certain more generic terms, such as 'feeling good', were mentioned discursively far more than they were specifically tagged, in this case by a factor of 175 times. Proportionally, other feelings such as 'amused' were tagged nearly as much as they were discussed in post text (0.81 times for every mention) possibly due to structural factors of the tagging feature, the corresponding emoticons, or simply the popularity of sharing such an emotion.

A basic temporal analysis of the data set, as we expected, revealed an increasing adoption of the feeling tag feature over time. Certain weekly distribution patterns were also observed in the larger discursive subset (Figure 5). As Figure 5 illustrates, feeling 'busy' was largest on Monday while feeling 'beautiful' was expressed significantly more on Fridays. Feeling 'drunk' increased before and during the weekend, while feeling 'frustrated' trended in the opposite direction, decreasing at the weekend.


Figure 5. Weekday Distribution of Selected Feeling Tags

Contextual Patterns: With Whom and Where?

Facebook affords its users the ability to add two extra tags when attaching a feeling tag to their post. Data on co-tagging with other Facebook friends can be used to investigate the feeling tags that are utilized most with other people (both in small and large groups). Data on cross-tagging with a location (be it a local place or a general city or region) can be used to analyze the 'where' dimension for the usage of feeling tags in status updates. Two examples are shown in Figures 6a and 6b below.

Collective Feelings. This research does not dive fully into investigating the social graph of actors with whom people share feelings. However, our feature extraction techniques can readily reveal whether feelings were expressed with others and how many people were included. Figure 6a illustrates an example post of a user that was 'feeling sad' together with friends. In total, users co-tagged other Facebook friends in *497,785* of the feeling-tagged posts (30.76%). One other person was added to feeling-tagged posts

8.83% of the time, while two people were added less often (2.36%). Tagging larger groups of three or more people accounted for 19.57% of all posts with a feeling tag.[10]

Figures 6a and 6b. Examples of Socialized and Localized Co-tags.

Localized Feelings. Figure 6b shows an example of when a status update includes a feeling and a location tag simultaneously. In this case footballer Lionel Messi is tagging the stadium of his football team, FC Barcelona (local tag instance). It would have also been possible for Messi to attach himself more broadly to Barcelona (region tag instance). Overall, less than 5% of posts were attached to a local place, and less than 2% were co-tagged with regional places. Situational emotions such as feeling 'tipsy', 'impatient' and 'drained' were the three feelings localized the most in public sharing, proportionally to how much they were discussed. These could be related to mobile uploads located at parties, shops, or gyms respectively, just to give some examples of user behavior that was observed. However in terms of sheer volume of co-tags, feeling excited was unsurprisingly the most common at locations overall.

Valence and Arousal of 'Facebook Feelings'

After investigating patterns in the use of feeling tags, we now turn to the assessment of the valence and arousal classifiers (trained by post text) that we built. Our two-dimensional classification of 'Facebook feelings' is shown in Figure 7. It includes 44 feelings that we selected for analysis from within the tagged posts. Both dimensions are depicted on a scale from -0.5 to +0.5. Feelings such as 'sleepy' and 'tired' are at the bottom of the graph's arousal axis with relatively neutral sentiment, for example. Of the 44 feelings, 26 overlap with those considered in prior research in emotion classification (Russell 1983; Scherer 2005) – we will compare our findings to these prior works in the next section. This means 18 feelings are unique to our dataset; some of these are particularly characteristic to social media and online slang ('pumped, 'awesome'), while others are informal variations of other feelings (e.g., 'pissed' is another way of expressing anger, frustration and annoyance in American English).

Our analyses indicate that the most commonly shared 'Facebook feelings' tend to be positive and characterized by high arousal (e.g., 'excited' and 'awesome'). The total volume of tags for just these two feelings is 177,640. This finding can be understood in light of the 'exhibitional' nature of Facebook (Hogan 2010), and is in line with prior findings on general online information sharing patterns of users (Berger and Milkman 2012; DeChoudhury, et al. 2012; Stieglitz and Dang-Xuan 2013). Feeling 'happy' and 'blessed' – mapped by our classifier as positive feelings with low to medium arousal levels – follow with a substantial 114,258 and 107,562 tags, respectively. On the negative side of the spectrum, feeling 'sad' is most commonly shared (70,835 tags), followed by 'annoyed' (16,838) and 'angry' (12,679). Sad and disappointed were detected as the feelings with the most negativity in their word usage. It is interesting to note that two of these high-volume feelings, 'awesome' and 'blessed' are specific to the Facebook's feeling tag feature option and have not been considered in prior research (e.g. Scherer 2005). The dictionary defines 'awesome' as inspiring great admiration (Oxford Dictionary). The typical meaning of feeling 'awesome' on Facebook, however, is more in line with the informal definition of 'awesome' as extremely good (ibid.); or cool (Internet Slang dictionary). In the world of Facebook, feeling 'awesome' allows people to 'produce an emotional solution' (Tucker and Ellis 2015) for a variety of positive and arousing contexts, in keeping with the curated, exhibitional nature of the medium (Hogan 2010). Feeling 'blessed', similarly, is an interesting adaptation of a term associated with a religious experience (having a sacred nature; connected with God) and using it to describe general gratitude for the good things happening in one's life.

Another noteworthy Facebook-specific feeling is 'meh' – characterized by relatively low arousal and neutral valence according to our data. According to Wikipedia, 'meh' is an interjection used as an

---

[10] This surprising level of social activity may be attributed to the nature of our public access to data. Our dataset may have a greater chance of capturing posts where larger groups allow for a higher probability that at least one member who re-shares the post has a 'public' setting for posts.

expression of indifference or boredom. It is often regarded as a verbal shrug of the shoulders and the use of the term demonstrates that the speaker is uninterested, or indifferent to the subject at hand. While 'meh' (like 'awesome') has not originated in Facebook, its semantic space is evolving due to the popularity of the medium ('meh' was tagged 6,278 times in our data set) that allows its users freedom in terms of the kinds of posts the tag is associated with.
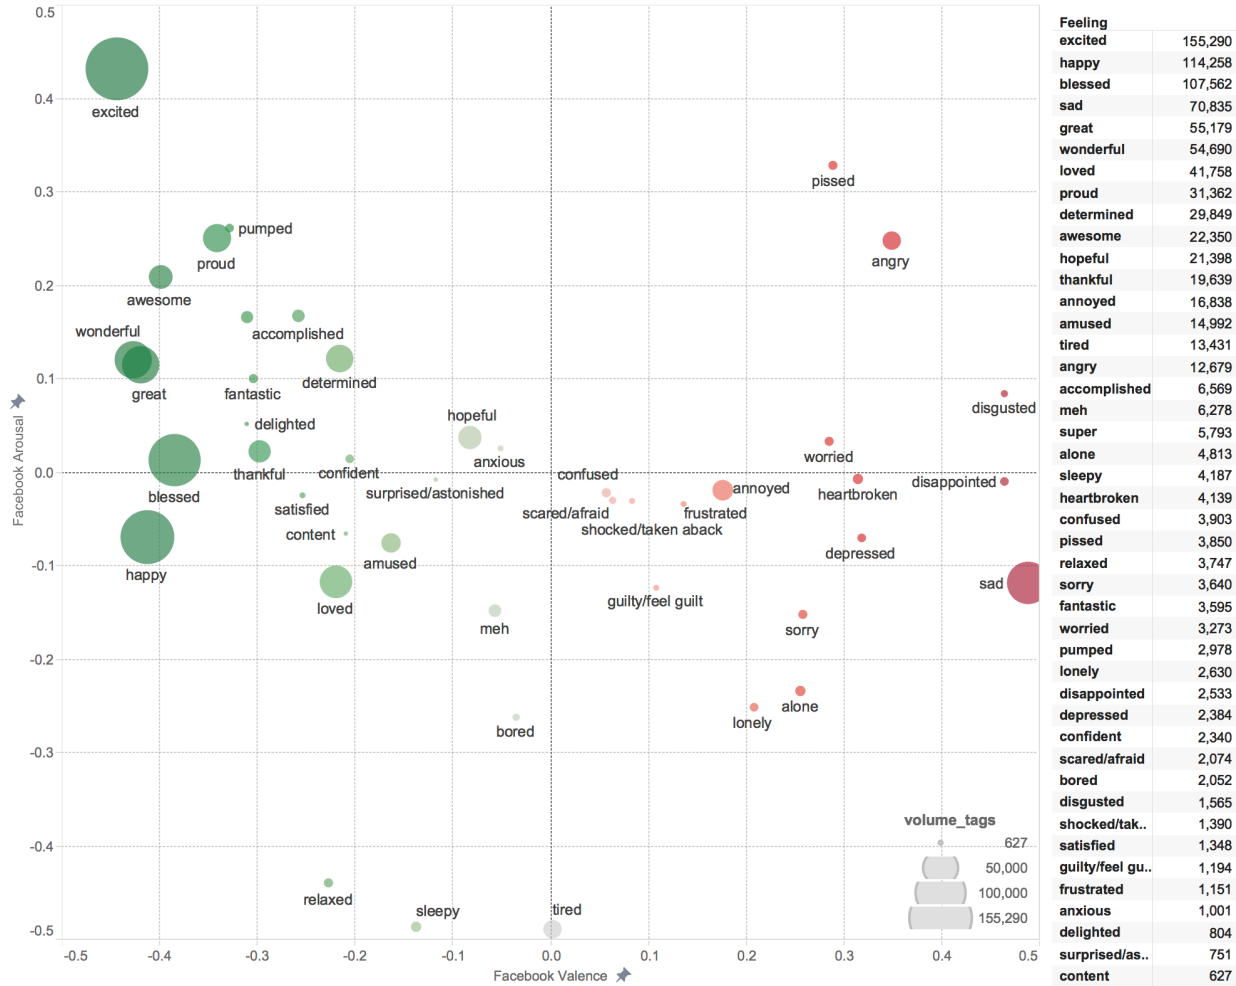

Feeling	
excited	155,290
happy	114,258
blessed	107,562
sad	70,835
great	55,179
wonderful	54,690
loved	41,758
proud	31,362
determined	29,849
awesome	22,350
hopeful	21,398
thankful	19,639
annoyed	16,838
amused	14,992
tired	13,431
angry	12,679
accomplished	6,569
meh	6,278
super	5,793
alone	4,813
sleepy	4,187
heartbroken	4,139
confused	3,903
pissed	3,850
relaxed	3,747
sorry	3,640
fantastic	3,595
worried	3,273
pumped	2,978
lonely	2,630
disappointed	2,533
depressed	2,384
confident	2,340
scared/afraid	2,074
bored	2,052
disgusted	1,565
shocked/tak..	1,390
satisfied	1,348
guilty/feel gu..	1,194
frustrated	1,151
anxious	1,001
delighted	804
surprised/as..	751
content	627

Figure 7. Classification of 'Facebook Feelings' (Post Volume Indicated by Circle Size) According to Valence and Arousal Dimensions

Comparing Semantic Spaces: Are Feelings Expressed on Facebook Different?

Of the 26 'Facebook feelings' in our study that overlap with those considered in prior research, 18 are common with Scherer's study (2005) and 16 with Russell's (1983). The comparison, including all three studies, is shown in Figure 8. Overall, our classification results are similar to those found in prior research (ibid.). The majority of comparable feelings fall into the same quadrants. Our classification of the 'tired' feeling based on the method of linguistic patterns (machine learning) from thousands of posts is exactly where Russell (1983) had previously positioned it based on qualitative research. Strong agreements with Scherer (2005) were observed for feelings such as 'angry', 'disappointed' and 'surprised' which all had less than 0.1 combined difference in valence and arousal positioning. The feelings that are unique to our classification fall into positions that make intuitive sense. For example, 'meh' is characterized by neutral valence and relatively low arousal, 'awesome' is characterized by medium level of arousal and positive valence. There are seven feelings that are common to all three studies: 'angry', 'sad', 'surprised', 'relaxed', 'excited', 'content', and 'bored'. If we look at the comparative positions of these seven feelings, we can observe some interesting patterns (Figures 9a & 9b) as discussed next.

Figure 8. Comparison of Findings: Russell (1983); Scherer (2005); Our Study

Extreme Feelings. Feeling 'angry' in our data is characterized by similar valence and arousal as in Scherer (2005). In both cases, anger is more negative but less aroused than in Russell (1983). In the Facebook context this may be explained by the existence of a high-arousal alternative – 'pissed'. Feeling 'sad' generally aligns with Russell's (1983) positioning, but Facebook's 'sad' is slightly more aroused and more negative. Our placement of 'excited' generally aligns with Russell's (1983) positioning but Facebook's 'excited' is slightly more aroused and more positive. In sum, for feelings that are generally considered strong or extreme (such as angry, sad and excited) extremeness seems to be amplified on Facebook in either both or one of the dimensions of valence and arousal. Further similar patterns could be observed in the case of feelings that were not covered by all three studies. For example, feeling 'disappointed' is characterized by higher negativity in the Facebook data than in the data by Scherer (2005). Interestingly, feeling 'happy', for example is characterized by a similar level of positivity as in Russell (1983), but with a lower level of arousal. It is possible that Facebook users prefer the alternative of feeling 'awesome', which indicates a similar level of positivity to 'happy', but a higher level of arousal.

Subtle Feelings. A rather different pattern emerges when analyzing feelings that are typically not considered strong or extreme (e.g., bored, relaxed, content). Facebook's feeling 'bored' emerges between Russell's (1983) and Scherer's (2005) 'bored' in terms of arousal but is more neutral in valence than either. Feeling 'relaxed' is characterized by both less arousal and less positivity than in prior findings (ibid.) and feeling 'content' is less positive. Similarly, Facebook's 'determined' is less positive but with a comparable level of arousal as in Scherer (ibid.), while Facebook's 'amused' is both less positive and less aroused than in Scherer (ibid.). For feelings that are generally considered mild (such as bored, relaxed, determined, amused), this mildness seems to be amplified on Facebook in either both or one of the dimensions so that the feelings become even more neutral.

Major Disagreements in Classification. The biggest discrepancy with prior theoretical classifications was with feeling 'anxious'. Numerically the two-dimensional classification revealed a distance of 0.41 for valence and 0.42 for arousal from Scherer's 2005 study. Scherer (2005) considers the feeling of 'anxious' to have both very high negativity and very low arousal. If one considers the term anxious to be a feeling in which the stimulus has yet to arrive, the valence classification by Scherer (2005) would make sense if it were almost certainly a negative stimulus. However, this is not necessarily the case, and many Facebook users seem to discuss being anxious in more neutral terms, even with a possibility of the anticipated event being positive. This is particularly true in case of major sports related events (we discuss this further in the organizational relevance subsection).
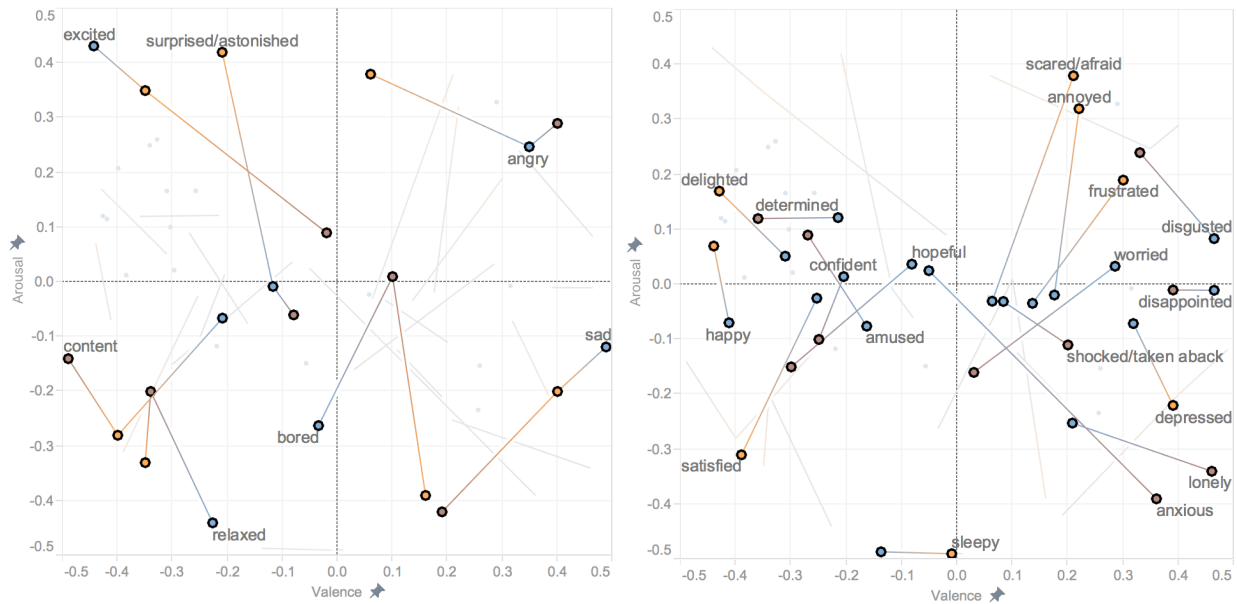
Figures 9a & 9b. Constellations of three-way and two-way comparisons in classification with past findings, where blue represents positions emerging from Facebook data.

Even with a few such significant disagreements, the data-driven placement of Facebook feelings has on average about 85% agreement with both Scherer's (2005) and Russell's (1983) classifications (average distances of .166 and .149 respectively). When comparing positions in the 3-way constellations above, it transpires that the positions of Facebook feelings differed less from Scherer (.187) or Russell (.160), than the results differed between the two past studies (average distance of .206). The higher degree of discrepancy with Scherer (2005) may stem from the fact that this study was conducted exclusively on German speakers while Russell's (1983) study included English speakers. This leads to further questions with regard to what aspects of 'Facebook feelings' may remain variant and invariant across different cultures, languages, and countries.

Summary of Findings

We take one of the first steps towards better defining user-categorized feelings in online social networks. Because of its feature for explicit tagging of posts with feelings and the popularity of the Facebook platform, we have been able to conduct our study on one of the most extensive and detailed data sets currently available. Some of the key patterns we observed, include: (1) feelings of excitement are the most widely shared, and positive-aroused feelings hold the most 'gravitational pull' in general, while there are few motivations to express neutrally-valenced feelings with moderate levels of arousal; (2) on the valence spectrum, the most negative feeling is that of sadness, greater than disappointment, anger or even disgust; (3) extreme and mild feelings tend to be exaggerated on Facebook; (4) the two-dimensional valence-arousal space of 'Facebook feelings' is qualitatively different from prior research (cf. Russell, 1983; Scherer, 2005); (5) yet variance between domain theorists (ibid.) is much higher than their individual variance with our empirical classification of 'Facebook feelings'.

Discussion

The facilitation of attaching feelings to online interactions via Facebook is of course very different to offline interactions. For one, the mechanisms on Facebook entice users via buttons, a dynamic menu of suggested feelings, and corresponding visual emoticons. Both the relative ordering of feeling tags in the vertical list and the attractiveness of corresponding emoticons may skew the popularity of different tags in use and, therefore, data sets such as ours. They also entice more users to volunteer feelings at large, helping increase the volumes of our feeling tag repository. One may also argue that the emoticon expression may also substitute facial expressions given their absence in online communication, therefore, being part of the feeling expression processes online. With this in mind, Facebook may indeed be providing users with not only a space for feeling and expressing emotions within (Ellis and Tucker 2015), but also a space where novel emotional scripts (Ashkanasy 2003) are created and performed. For example, feeling 'meh' arguably fills a void of where very few feelings of indifference have been documented in past research; 'meh' also helps people express such feelings of disinterest in the Facebook medium in ways that other traditionally socialized feelings (e.g., 'bored') do not.

From a social psychological perspective, our results also indicate that there is a possibility for re-alignment of certain emotions (for example, anxious, excited, sad, angry) that may be interpreted and utilized differently by millions of Facebook users. As social media become increasingly integrated into lived experiences as the "technologies of the self" (Foucault et al. 1988), the emergence and establishment of new forms of online expression of subjective feelings and the performance of the social self becomes co-determined by social and technical aspects resulting in novel phenomenological modes of "technological intersubjectivity" (Vatrapu and Suthers 2009). For example, as the recent controversial study of emotional contagion on Facebook (Kramer 2014) showed, algorithmic manipulation of users' subjective feelings is technically possible even if such practices are not ethically justifiable and/or legally permissible. Our findings with respect to the contextual use of the feeling tag "anxious" (and the use of co-tagging of other users and places) might be relevant for combining appraisal theories of emotional processes widely employed in the computational modeling of emotions of virtual agents (Marsella & Gratch 2014; Gratch & Marsella 2001) with the dimensional approach widely adopted in social psychology and information systems that we employed in this paper. We believe that big social datasets of human emotions and feelings can support much needed cross-disciplinary research that investigates causes, categories and consequences of the socio-technical phenomena. In that regard, understanding the interactional processes of subjective feeling expressions in addition to the linguistic outcomes in terms of the postings made would lead to better modeling of emotions for business applications such as conversational agents, virtual store assistants, customer avatars etc. (cf. Marsella and Gratch 2014; Gratch and Marsella 2001). In addition, our findings contribute to the growing literature on investigating emotions in big social data termed "hedonometrics" (Dodds et al. 2011). Our results show that the temporal patterns of subjective feelings on Facebook are aligned with prior findings with Twitter datasets. Future analysis of the co-tagging of Facebook feelings with other Facebook users will allow us to compare the expressions of subjective feelings on other social media (Bliss et al. 2012). This helps us better understand the conceptual space of subjective feelings on not only individual social media platforms such as Facebook and Twitter but also in online media in general.

Implications for Practice: Automatic Feelings Meter

Opinions from the crowd are increasingly being reported with the help of social media and business intelligence tools utilizing automatic sentiment analysis. However, taking the temperature of text on aggregate does not provide any indication of which granular feelings may be responsible for overal swings in positive or negative mood levels from the crowd. Furthermore, most tools are still one-dimensional by only focusing on valence (Talkwalker, FanpageKarma, Topsy). For example the two feelings – sad and angry – may be at similar levels on the valence (sentiment) spectrum, yet their arousal levels are notably different (Figure 8). With increasingly high volumes of conversations across social media, the rapid detection of certain core feelings may be of significant strategic value to industry practitioners monitoring product launches, campaigns and public relations milestones. For these reasons, we have developed a research prototype, a Feelings Meter, to leverage our trained classifiers towards practical applications.

It is common in emotions research to divide discrete feelings into a smaller number of 'core' categories, such as joy, anger, sadness, fear and excitement (Ekman 1992; Parrott 2001). From a practical perspective, such abstraction may be useful when building a social media analytics tool focusing on affect

(cf. Abbasi and Chen, 2008; Stieglitz and Dang-Xuan, 2013). Most organizations will see value in understanding how their customers feel and which feelings they express on the Facebook wall of the organization. Such insights could be used for improvements in marketing strategy, community engagement, customer service, reputation management, and the discovery of new business opportunities (Goh, et al., 2013; Holsapple, et al., 2014; Kurniawati, et al., 2013). Accordingly, in an initial test of the practical usefulness of our data, we have built a five-way classifier that would enable organizations to do just that - feed their Facebook wall data into a Feelings Meter that, in turn, detects feeling groups prevalent on their wall as output (Figure 10).


Figure 10. Feelings Meter Demo Version

At this stage, we have only constructed a simple test version of the Feelings Meter (cssl.cbs.dk/software/feelingsmeter). The user can enter text that is then evaluated for five feeling groups (angry, animated, empowered, fearful and joyous) on a scale of 0 to 1.0. Each feeling group leverages a subset of discrete feelings that are traditionally placed in the same group (e.g., joy consists of feeling happy, wonderful, delighted, etc.). In future work, we will progress towards ways of systematically assessing accuracy by eliciting human judgments about the level of different feelings in test data, and then compare these with the output. The test version of the Feelings Meter itself has limitations. Since the current version assumes *all* texts express feelings, so it is unable to account for and exclude neutral texts. As a result, it will always attempt to classify a text according to the five 'core' feeling groups. Thus, if a text does not have any emotional gravity, the result will be meaningless. While we have not done systematic assessment of the Feelings Meter output, we have noted an apparent bias towards the groups joyous and animated. This is likely due to the nature of Facebook posts, which tend to express positive, high arousal feelings more frequently (see Figure 7).

Organizational Relevance

Organizations could potentially use this detection of feelings for both brand positioning inside their community and conversation monitoring outside their community. By tracking user conversations in relation to the brand, businesses can not only detect specific feeling groups with the help of the Feelings Meter, but also analyze the valence and arousal levels of the conversation. Figure 11 illustrates an analysis of conversations about the FIFA brand name during the corruption scandal in May 2015 (The Guardian 2015).


Figure 11. FIFA Brand Conversation

When we test the Feelings Meter on the public conversation about FIFA on Facebook (Figure 11) we see several changes in emotional tone at the point of the scandal. First, arousal level of the conversation remains relatively unchanged (1), however a significant spike in collective negativity is strongly apparent

(2). It is likely that football attracts high arousal conversations at all times, whereas the spike in negativity is clearly telling of the negative public reaction to the scandal. Second, even more dramatic variations occur when contrasting two discrete emotions: anger and joy. Fluctuations in joyous levels seem to stabilize somewhat after the event (3), while a pronounced surge in anger occurs immediately after FIFA executives were arrested in Switzerland (4).

Two separate organizational application areas can be identified for the online marketing of a brand such as FIFA: one within the owned channel (brand Facebook page and its fans) and the other towards the earned conversation by the crowd (those who publicly mention the brand on Facebook, Twitter, etc.).

Brand Positioning and Emotional Alignment. A social media manager can compare the emotional tone of the conversation before and after key postings by the organization, as well as detect an *'emotional alignment'* within the subsequent comment chain that may be generated in reaction to the published post. This internal organizational application of the Feelings Meter tool seeks to inform practitioners of whether there is alignment between the purposeful publishing of emotional posts (e.g., happy or excited) and the corresponding reactions received. The tool crucially could measure which emotions are detected from the fans who respond to the post. It may further be applied towards goals such as production of content that goes viral, or desired emotional conversational output from the brand community.

Conversation Monitoring and Detection from the Crowd. A more basic application is that of listening to the crowd at large, a radar of sorts, monitoring what marketers refer to as 'earned media'. This includes all brand mentions in user-generated content online. This public chatter on social media happens at any time and is typically beyond any strategic content or initiation by the brand itself.

Concluding Remarks

The volunteered feeling declarations in combination with corresponding text provide a unique opportunity to leverage machine learning on a training set far more appropriate than one-dimensional movie and product ratings traditionally used to train text classifiers for positive, negative and neutral predictions of text sentiment. In this study we take one of the first steps towards better defining user-volunteered feelings in online social networks. Because of this tagging system and its high adoption rate, the Facebook platform currently has the most comprehensive data set for doing so. Business organizations can benefit from more informative classifications, which in turn can be leveraged for better monitoring of the large conversation streams that revolve around brands online, while empowering decision-makers to take action in their online communities.

A notable strength of this study relates to its unique data set that allows us to leverage big social data to shed light on the emotional lives of the public (typically a topic investigated with small data studies). Furthermore, the nature of the data also allows us to draw on the folksonomic wisdom of the crowds. Past classifications of emotions reflect the opinions of small numbers of study participants as well as those of the researchers (cf. Scherer 2005). In Facebook, while the list of feeling tags is pre-defined, users are free to choose from this list and also to create their own. The sheer volume of data thus allows us to leverage the contributions of the English speaking population who volunteer feeling tags as they see appropriate. The spread of user-appropriated tags is likely to influence how others use them across time and space. Traditional survey studies are performed at specific times and spaces, not allowing subjects to appropriate emotions in an embedded fashion at any and every point in time in their daily lives. Our data-driven approach from big social data has allowed the patterns in feeling tag use to emerge from millions of posts, letting the data speak for itself, and to reveal observable differences from past assumptions.

Acknowledgements

References

Abbasi, A., and Chen, H. 2008. "CyberGate: a design framework and system for text analysis of computer-mediated communication," *MIS Quarterly* (32:4), pp. 811-837.

Abrahams, A. S., Jiao, J., Fan, W., Wang, G. A., and Zhang, Z. 2013. "What's buzzing in the blizzard of buzz? Automotive component isolation in social media postings," *Decision Support Systems* (55:4), pp. 871-882.

Ashkanasy, N. M. 2003. "Emotions in organizations: A multilevel perspective," in *Multi-Level Issues in Organizational Behavior and Strategy (Research in Multi Level Issues, Volume 2)*, F. Dansereau and F. J. Yammarino (ed.) Emerald Group Publishing Limited, pp.9 – 54.

Asur, S. and Huberman, B. A. 2010. "Predicting the Future with Social Media", In proceedings of *International Conference on Web Intelligence and Intelligent Agent Technology (WI-IAT) 2010 (1)* IEEE, Toronto, ON, pp.492–499.

Bagozzi, R.P. 2007. "The Legacy of the Technology Acceptance Model and a Proposal for a Paradigm Shift," *Journal of the Association for Information Systems* (8:4), pp. 244–254.

Barnaghi, P., Ghaffari, P. and Breslin, J. 2015. Text Analysis and Sentiment Polarity on FIFA World Cup 2014 Tweets. Available at: http://www.large-scale-sports-analytics.org/Large-Scale-Sports-Analytics/Submissions_files/paperID11-KDD'15(camera-ready)-Peiman.pdf [Accessed September 4, 2015].

BBC News. 2015. Facebook has a billion users in a single day, says Mark Zuckerberg. *BBC News - Technology*. Available at: http://www.bbc.com/news/world-us-canada-34082393 [Accessed September 5, 2015].

Bliss, C., Kloumann, I. and Harris, K. 2012. "Twitter reciprocal reply networks exhibit assortativity with respect to happiness," *Journal of Computational Science* (*3:5)*, pp. 388-397.

Barsade, S. G., and Gibson, D. E. 2007. "Why does affect matter in organizations?" *The Academy of Management Perspectives* (*21*:1), pp. 36-59.

Beaudry, B.A., and Pinsonneault, A. 2010. "The Other Side of Acceptance: Studying the Direct and Indirect Effects of Emotions on Information Technology Use," *MIS Quarterly* (34:4), pp. 689-710.

Benozzo, A., and Colley, H. 2012. "Guest Editorial, Emotion and learning in the workplace: critical perspectives," *Journal of workplace learning* (24:5), pp. 304-316.

Chau, M., and Xu, J. 2012. "Business intelligence in blogs: Understanding consumer interactions and communities," *MIS Quarterly* (36:4), pp. 1189-1216.

Culnan, M. J., McHugh, P. J., and Zubillaga, J. I. 2010. "How large US companies can use Twitter and other social media to gain business value," *MIS Quarterly Executive* (9:4), pp. 243-259.

De Guinea, A. O., and Markus, M. L. 2009. "Why break the habit of a lifetime? Rethinking the roles of intention, habit, and emotion in continuing information technology use," *MIS Quarterly* (33:3), pp. 433-444.

Dodds, P. S., Harris, K. D., Kloumann, I. M., Bliss, C. A., and Danforth, C. M. 2011. "Temporal patterns of happiness and information in a global social network: Hedonometrics and Twitter." *PloS one, 6*(12), e26752.

Ekman, P. 1992. "An argument for basic emotions," *Cognition & emotion* (*6:3-4)*, pp. 169-200.

Ellis, D., and Tucker, I. 2015. *Social Psychology of Emotion*. Thousand Oaks, CA: SAGE Publications Inc.

Facebook Company Info - Statistics. Available at: http://newsroom.fb.com/company-info/ [Accessed May 2, 2015].

Ferrara, E., Varol, O., Davis, C., Menczer, F., and Flammini, A. 2014. "The rise of social bots," *arXiv preprint arXiv:1407.5225*.

Fineman, S. 2008. The Emotional Organization: Passions and Power, Oxford, UK: Blackwell.

Foucault, M., Martin, L. H., Gutman, H., & Hutton, P. H. 1988. Technologies of the self: A seminar with Michel Foucault Univ of Massachusetts Press.

Goh, K.-Y., Heng, C. S., and Lin, Z. 2013. "Social Media Brand Community and Consumer Behavior: Quantifying the Relative Impact of User-and Marketer-Generated Content," Information Systems Research (24:1), pp. 88-107.

Grant, A.M. 2013. "Rocking the boat but keeping it steady: the role of emotion regulation in employee voice," *Academy of Management Journal* (56:6), pp. 1703-1723.

Gratch, J. and Marsella, S. 2001. "Tears and fears: Modeling emotions and emotional behaviors in synthetic agents," in *Proceedings of the fifth international conference on Autonomous agents*, Montreal, Canada, pp. 278-285.

Hogan, B. 2010. "The presentation of self in the age of social media: Distinguishing performances and exhibitions online," *Bulletin of Science, Technology and Society* (30:6), pp. 377-386.

Holsapple, C., Hsiao, S. H., and Pakath, R. 2014. "Business Social Media Analytics: Definition, Benefits, and Challenges," in *Proceedings of the Americas Conference on Information Systems 2014*, Savannah, GA.

Kamvar, S.D. and Harris, J. 2011. "We feel fine and searching the emotional web," in *Proceedings of the fourth ACM international conference on Web search and data mining - WSDM '11*, Kowloon, Hong Kong, p.117.

Kopelman, S., Rosette, A.S., and Thompson, L. 2006. "The three faces of Eve: Strategic displays of positive, negative, and neutral emotions in negotiations," *Organizational Behavior and Human Decision Processes* (99:1), pp. 81-101.

Kramer, A. D., Guillory, J. E., and Hancock, J. T. 2014. "Experimental evidence of massive-scale emotional contagion through social networks," in *Proceedings of the National Academy of Sciences* (111:24), pp. 8788-8790.

Kurniawati, K., Shanks, G., and Bekmamedova, N. 2013. "The Business Impact of Social Media Analytics," in *Proceedings of the 21th European Conference on Information System,* Utrecht, The Netherlands, p. 48-61.

Lawler, E.J. and Thye, S.R. 1999. "Bringing Emotions into Social Exchange Theory," *The Annual Review of Sociology* (25), pp. 217-244.

Lindebaum, D., and Cartwright, S. 2010. "A Critical Examination of the Relationship between Emotional Intelligence and Transformational Leadership," *Journal of Management Studies* (47:7), pp. 1317-1342.

Mandviwalla, M., and Watson, R. 2014. "Generating Capital from Social Media," *MIS Quarterly Executive* (13:2), pp. 97-113.

Marsella, S.C. and Gratch, J. 2014. "Computationally modeling human emotion," *Communications of the ACM* (*57:*12), pp.56–67.

McGrath, K. 2006. "Affection not affliction: The role of emotions in information systems and organizational change," *Information and Organization* (16:4), pp. 277-303.

Mukkamala, R. R., Sørensen, J. I., Hussain, A., and Vatrapu, R. 2015. "Social Set Analysis of Corporate Social Media Crises on Facebook," *Proceedings of the19th IEEE International Conference on Enterprise Distributed Object Computing (IEEE EDOC 2015).*

Pang, B., and Lillian L. 2008. "Opinion mining and sentiment analysis," *Foundations and trends in information retrieval* (2:1-2), pp. 1-135.

Pang, B., Lillian L., and Shivakumar V. 2002. "Thumbs up?: sentiment classification using machine learning techniques." in *Proceedings of the ACL-02 conference on Empirical methods in natural language processing-Volume 10* Association for Computational Linguistics, Philadelphia, PA, pp. 79-86.

Parrott, W. 2001. *Emotions in social psychology: Essential readings*. Philadelphia, Chicago: Psychology Press.

Pew Social Media Report 2015. Available at: http://www.pewinternet.org (January 2015). [Accessed May 2, 2015].

Russell, J. A. 1983. "Pancultural aspects of the human conceptual organization of emotions," *Journal of Personality and Social Psychology* (45:6), pp. 1281 – 1288.

Russell, J. A. 2003. "Core affect and the psychological construction of emotion," *Psychological review* (110:1), pp. 145-172.

Scherer, K. R. 2005. "What are emotions? And how can they be measured?" *Social science information* (44:4), pp. 695-729.

Standage, T. 2013. Writing on the Wall: Social Media-the First 2,000 Years, New York, NY: Bloomsbury.

Stieglitz, S., and Dang-Xuan, L. 2013. "Emotions and information diffusion in social media—Sentiment of microblogs and sharing behavior," *Journal of Management Information Systems* (29:4), pp. 217-248.

Stein, M. K., Newell, S., Wagner, E. L., and Galliers, R. D. 2014. "Felt quality of sociomaterial relations: Introducing emotions into sociomaterial theorizing," *Information and Organization* (24:3), pp. 156-175.

Stein, M. K., Newell, S., Wagner, E.L., and Galliers, R. 2015. "Coping with Information Technology: Mixed Emotions, Vacillation, and Nonconforming Use Patterns," *MIS Quarterly* (39:2), pp. 367-392.

Tableau 9.0 - Business Intelligence and Analytics Software. Available at: http://www.tableau.com/ [Accessed April 23, 2015].

The Guardian. 2015. *Fifa officials arrested on corruption charges as World Cup inquiry launched* Available at: http://www.theguardian.com/football/2015/may/27/several-top-fifa-officials-arrested [Accessed September 5, 2015].

Thompson, M. 2012. "People, practice, and technology: Restoring Giddens' broader philosophy to the study of information systems," *Information and Organization* (22:3), pp. 188-207.

Tufekci, Z. 2014. "Big Questions for Social Media Big Data: Representativeness, Validity and Other Methodological Pitfalls," in *ICWSM'14: Proceedings of the 8th International AAAI Conference on Weblogs and Social Media*, Ann Arbor, MI, p.10.

Vatrapu, R. 2013. Understanding social business. in *Emerging Dimensions of Technology Management*. India: Springer, pp. 147-158.

Zimmerman, C., Chen, Y., Hardt, D., and Vatrapu, R. 2014. "*Marius, the giraffe: a comparative informatics case study of linguistic features of the social media discourse*," Paper presented at the *Proceedings of the 5th ACM international conference on Collaboration across boundaries: culture, distance and technology*, Kyoto, Japan.