Geography of Emotion: Where in a City are People Happier?

Luciano Gallegos
USC Information Science Institute
Marina del Rey, CA 90292
luciano.gallegos@gmail.com

Kristina Lerman
USC Information Science Institute
Marina del Rey, CA 90292
lerman@isi.edu

Arthur Huang
Tarleton State University
Stephenville, TX 76402
ahuang@tarleton.edu

David Garcia
ETH Zurich
8092 Zurich, Switzerland
dgarcia@ethz.ch

ABSTRACT

During the last years, researchers explored the geographic and environmental factors that affect happiness. More recently, location-sharing services provided by the social media has given an unprecedented access to geo-located data for studying the interplay between these factors on a much bigger scale. Do location-sharing services help in turn at distinguishing emotions in places within a city? Which aspects contribute better at understanding happier places? To answer these questions, we use data from Foursquare location-sharing service to identify areas within a major US metropolitan area with many check-ins, i.e., areas that people like to use. We then use data from the Twitter microblogging platform to analyze the properties of these areas. Specifically, we have extracted a large corpus of geo-tagged messages, called tweets, from a major metropolitan area and linked them US Census data through their locations. This allows us to measure the sentiment expressed in tweets that are posted from a specific area, and also use that area's demographic properties in analysis. Our results reveal that areas with many check-ins are different from other areas within the metropolitan region. In particular, these areas have happier tweets, which also encourage people living in it or from other areas to commute longer distances to these places. These findings shed light on the influence certain places play within a city regarding people's emotions and mobility, which in turn can be used for city planners for designing happier and more equitable cities.

Keywords

location-sharing services; mobility; sentiment analysis; social media

WWW'16 Companion, April 11–15, 2016, Montréal, Québec, Canada.
ACM 978-1-4503-4144-8/16/04.
http://dx.doi.org/10.1145/2872518.2890084 .

1. INTRODUCTION

Along the years, psychologists and economists designed surveys to quantify an individual's subjective level of happiness, subjective well-being, or satisfaction with life, which then allowed them to study how responses to survey questions are related to socioeconomic and demographic factors. More recently, researchers explored the geographic and environmental factors that affect happiness. They found that proximity to amenities, such as the coast or major routes of transportation, lead to higher levels of subjective well-being, while proximity to a landfill negatively affects well-being [1]. Such findings offer guidelines for city planners and policy makers for designing urban areas that promote happiness and maximize equity in the distribution of resources.

The rise of social media — and location-sharing services in particular — has given researchers an unprecedented access to geo-located data for studying the interplay between geography and happiness on a much bigger scale and more precise temporal and geographical resolutions. Despite these advantages, social media data sources suffer from self-selection and demographic biases [33], and thus provide an alternative approach to survey data rather than a strict improvement. It is in the combination of different methodologies that we can derive new knowledge, rather than by using each separately and arguing about which one is best.

The availability of text on social media platforms also enabled researchers to analyze the sentiment of the messages and the emotional state of individuals posting them. Sentiment analysis has received much attention from the research community [20], since it allows people to monitor sentiment on a global scale [16] — an impossible task if one had to rely on surveys to measure people's emotional states [27]. For example, Kramer [21] built a sentiment score for Facebook status updates and found that it correlates well with the self-reported satisfaction with life at the national (macro) level. Eagle et al. [8] showed that the subjective well-being of communities strongly correlates with network diversity, where members of well-off communities have diverse networks while members of economically and socially disadvantaged communities have insular social relations. Alshamsi et al. [3] studied the effectiveness of social media in mapping happiness at finer spatial resolution and found that happy areas tend to interact with other happy areas (i.e., homophily), although they did not use other urban indicators such as demographics or mobility [5, 24].

Do location-sharing services help us at distinguishing emotions in places within a city? Which aspects contribute better at understanding happier places? Aiming at answer these questions and using previous methods [18, 23, 27, 28, 34], in this paper we combined social media data from Twitter and Foursquare, with demographic data from the US Census to carry out a deeper fine-grained analysis of geography and emotion within a city. Specifically, we used the Foursquare location-sharing service (the so-called "check-ins") to identify areas within a major US metropolitan area. These are the places with amenities, such as parks, restaurants, public transportation, and gyms. We then analyzed properties of these places by looking at geo-tagged messages, or tweets, coming from those areas, and linked the tweets to US Census tracts through their locations. This allows us to link the sentiment expressed in tweets that are posted from different census tracts with those tracts' demographic properties.

Our results reveal that areas with more check-ins are different from other areas within the city. In particular, these areas have happier tweets, which also encourage people living in it or from other areas to commute longer distances to these places. These findings shed light on the influence certain places play within a city regarding people's emotions and mobility, which in turn can be used by city planners to design happier and more equitable cities.

2. RELATED WORKS

Research works reported [6, 24] that Foursquare users usually check-in at venues they perceived as more interesting and express actions similar to other social media, such as Facebook and Twitter. Foursquare check-ins are, in many cases, biased: while some users provide important feedback by checking-in at venues and share their engagement, others subvert the rules by deliberately creating unofficial duplicate and nonexistent venues [7]. The high availability of Foursquare and Twitter data transmitted from mobile devices has also been subject to human mobility research [9]. More specifically, some researchers used Radius of Gyration (r_g) [18] to characterize and quantify human mobility. For example, Noulas et al. [24] applied r_g to conduct a large-scale study of user behaviour and Foursquare check-ins with 700K users spanning a period of more than 100 days. The study revealed users' temporal and mobility patterns (the majority of users moved between 1 and 10 km and expended 100 and 2000 minutes to do so) in urban locations. Usually, human mobility is measured at the individual's level of granularity, disclosing the users' profile as well as their personal mobility patterns, which potentially discloses information that the user may prefer to keep private [3].

Mitchell et al. [23] generated taxonomies of US states and cities based on their similarities in word use and estimates the happiness levels of these states and cities. Then, the authors correlated highly-resolved demographic characteristics with happiness levels and connected word choice and message length with urban characteristics such as education levels and obesity rates, showing that social media may potentially be used to estimate real-time levels and changes in population-scale measures, such as obesity rates. Eagle et al. [8] built communication networks from phone records across the entire United Kingdom, cross-referenced it with Census data, and showed that members of well-off communities have diverse networks, while members of economically and socially disadvantaged communities have insular social relations. Quercia et al. [28, 34] used the Index of Multiple Deprivation (i.e., qualitative study of deprived areas in the UK local councils) to compute happiness based on small areas, providing promising fine-grained, micro-level results. Alshamsi et al. [3] studied the effectiveness of social media in mapping happiness with finer spatial resolution and, similar to [27], found that happy areas tend to interact with other happy areas (i.e., homophily), although other indicators such as demographic data and human mobility were not used in their research [5, 24].

3. DATA

We collected a large body of tweets from Los Angeles County over the course of 4 months, starting in July 2014. First, we used Twitter's location search API to collect tweets from an area that included Los Angeles County. We then used Twitter4J API to collect all (timeline) tweets from users who tweeted from within this area during this time period. A portion of these tweets were geo-tagged, i.e. tweets with geographic coordinates attached to them. In all, we collected 6M tweets, of which 700K made by 24K distinct users were geo-tagged.

We localized geo-tagged tweets to tracts from the 2012 US Census[1]. A *tract* is a geographic region defined for the purpose of taking a census of a population, containing about 4,000 residents on average, and is designed to be relatively homogeneous with respect to socioeconomic characteristics of that population. We included only Los Angeles County tracts in the analysis.

Some Foursquare users link their accounts to Twitter, so that their check-ins will be visible to their Twitter followers. Such check-in tweets were automatically generated and had a specific format, e.g., "I'm at 1K Studios (Burbank, CA) $http://t.co/3W5ymDM5EI$", "I'm at @Specialtys Cafe & Bakery in Glendale, CA $https://t.co/IeHOY6Bbbz$", "I'm at Bossa Nova Brazilian Cuisine - @bossanovaeats (West Hollywood, CA) $http://t.co/pGHsMVGE3v$". We created parsers to extract the location and venue from these tweets. In all, we extracted 5,863 check-ins from 687 tracts around Los Angeles County. The Los Angeles tract map in Figure 1 are coloured by the number of check-ins within their boundaries, showing that just a fraction of the collected data is geo-tagged.

4. METHODS

The field of sentiment analysis [26] aims at developing tools that process text to quantify subjective states, including opinions and emotions. Two recent independent surveys evaluated different sentiment analysis tools in various social media [17] and in a benchmark of datasets from Twitter [2]. Across social media, one of the best performing tools is SentiStrength [31], which also was shown to be the best unsupervised tool for tweets in various contexts [2]. SentiStrength quantifies the emotions expressed in text by applying a lexicon and taking into account intensifiers, negations, misspellings, idioms, and emoticons. We apply the standard English version of SentiStrength to each tweet in our dataset, quantifying Positive and Negative sentiment in

[1]American Fact Finder (http://factfinder.census.gov)

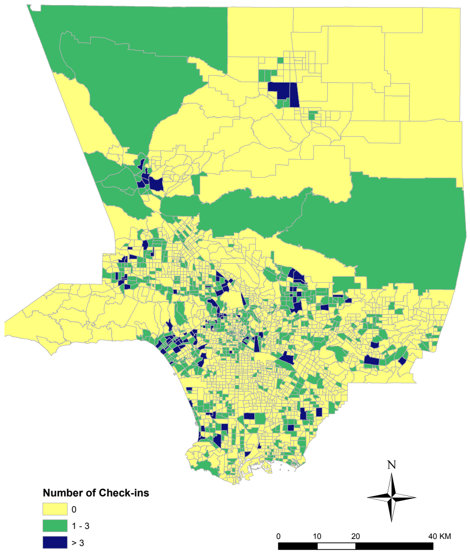

Figure 1: Los Angeles Census tract map. Tracts are colored by the number of Foursquare check-ins within them.

a way that is consistent with the Positive and Negative Affect Schedule (PANAS) in psychology [36]. Beyond its accuracy, SentiStrength has been shown to perform very closely to human raters in validity tests [31] and has been applied to measure emotions in product reviews [15], online chatrooms [12], Yahoo answers [22], and Youtube comments [14].

Research in psychology shows that emotional experiences contain components in more than two dimensions [11], calling for extended analysis that includes multidimensional aspects of emotions. When measured through text, emotional meanings can be quantified through the application of the semantic differential [25], a dimensional approach that quantifies emotional meaning in terms of Valence, Arousal, and Dominance [30]. The dimension of Valence quantifies the level of pleasure expressed by a word, Arousal measures the level of activity induced by the emotions associated with a word, and Dominance quantifies the level of subjective power experienced during an emotion. The state of the art in the quantification of these three dimensions is the lexicon of Warriner, Kuperman, and Brysbaert (WKB) [35], which includes scores in the three dimensions for almost 14,000 English lemmas. To quantify the Valence, Arousal, and Dominance expressed in a tweet, we lemmatize its content and apply the lexicon to compute mean values of the three dimensions as in [19]. Thanks to the size of the lexicon, we find emotional terms in 82.39% of the tweets in our dataset, producing a multidimensional measure of emotion aggregates as expressed through tweets.

Mobility patterns are well-correlated with demographics and individual's socioeconomic status [5]. Studies of human mobility usually focus on either the small scale (e.g., travel modes of individuals' daily commutes) or the large scale (e.g., air-travel patterns to track the spread of epidemics over time). Some researchers adapted concepts from physics, such as the *radius of gyration* (r_g), to characterize human mobility [18]. The r_g is the standard deviation of distances between the user's locations (given by geo-tagged tweets) and the center of mass of those locations, and measures both how frequently and how far a user moves. A low r_g indicates a user who travels mainly locally (with tweets mainly concentrated in a small geographic area), while a high radius of gyration indicates a user whose tweets are spread far apart spatially. The r_g for a user is defined as [18]:

$$r_g = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(r_i - r_{cm})^2} \quad (1)$$

where n is the number of geo-tagged tweets posted by that user, and $(r_i - r_{cm})$ is the distance between a particular tweet r_i and the user's center of mass r_{cm}. The latter is simply the average location of all tweets. In this paper, we apply the r_g to quantify a Twitter user's mobility.

5. RESULTS

The emotions expressed in tweets posted from different places around Los Angeles County is first analysed. We find that geo-tagged tweets related to areas with many check-ins have happier emotions, and are likely to have higher mobility. We also extend our analysis to explore how demographic factors contribute to these observations.

Sentiment analysis

We use WKB and SentiStrength scores, described in Section 4, to measure the emotional content of geo-tagged tweets from the Los Angeles tracts. Similarly to [13], we set the score scale for Valence, Arousal and Dominance ranging from 1 (low) to 9 (high) with 5 being the neutral score, and the scale for Positive and Negative ranging from 1 (low) to 5 (high). The mean of the scores is calculated for each tract and is in turn organized into groups of check-ins within a certain empirical threshold. The mean of each group is calculated and results are compared across groups.

The Figure 2 depicts the mean of sentiment scores with a 95% confidence interval and is organized in three groups: the original 1718 tracts with geo-tagged tweets as "$check-ins \geqslant 0$", 687 tracts (approximately 40% of the original tracts) as "$check-ins \geqslant 1$", and 162 tracts (9%) as "$check-ins \geqslant 3$". Due to the lack of sufficient data for thresholds above $check-ins \geqslant 3$, the following analysis is restricted to these three groups. Results show that Valence, Arousal and Dominance WBK, as well as Positive SentiStrength show higher mean sentiment scores in groups containing more geo-tagged tweets with check-ins. Regarding the Negative SentiStrength, the more geo-tagged tweets with check-ins a group has the less negative written messages are. It is likely that people tweeting from these tracts express happier and less negative emotions due to their moment of experience, e.g., having a nice leisure time by the beach.

Are these check-in groups different indeed from each other in terms of mean scores? To answer this question, and since the distribution shapes are non-Gaussian, we further compute the *Wilcoxon signed-rank test* — a non-parametric statistical hypothesis test used when comparing two related samples. For this purpose, we include in each group the standard deviation of each mean sentiment score. Then, we test the null hypothesis that no difference exists when comparing mean emotion scores across groups. The test is performed *two-by-two and in acceding order*, so that we test the groups $check-ins \geqslant 0$ and $check-ins \geqslant 1$, and the groups $check-ins \geqslant 1$ and $check-ins \geqslant 3$.

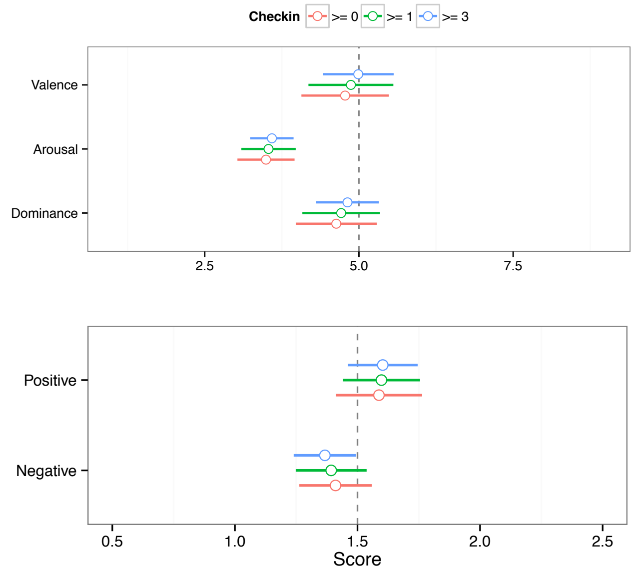

Figure 2: Means and 95% confidence interval from emotion scores into group of tracts containing check-ins — WBK lexicon (top) and SentiStrength (botton). Higher means of emotion scores of geo-tagged tweets are related to groups containing more quantity of check-ins (plotted in green and blue).

The results described in the Table 1 provide strong evidences that the mean sentiment scores between groups are different. For instance, $check-ins \geqslant 0$ has lower Valence mean score compared to $check-ins \geqslant 1$ at a high significance level ($p < 0.001$), leading us to reject the null hypothesis that Valence scores are similar — and that groups containing more tweets with check-ins have happier texts.

Why are tweets from tracts with many check-ins more emotional? To get insight, we look at the words that are commonly used in these tweets. Figure 3 shows the world cloud of tweets from a single tract without check-ins and one tract that has many check-ins. The latter tract has words like “beach” and “playa” (Spanish for “beach”), in addition to “paseo” and “esplanade”, which suggest pleasant places to stroll. While tweets from the first tract have “Shrine Auditorium”, which is a popular venue for concerts, they have fewer words associated with pleasant experiences, such as going to the beach, or strolling with friends. Though deeper analysis is required, these results suggest again that place with check-ins offer pleasant amenities, such as the beach, that attract people to those areas.

Human mobility

Inspired by works in urban mobility [5, 18] as detailed in Section 4, we compute the “radius of gyration” (r_g) that represents the average distance between the user’s locations and the tract’s center of mass. We aggregate these values to quantify tract mobility as the “average radius of gyration”' (R_g) of all users tweeting from that tract:

$$R_g = \frac{1}{N}\sum_{i=1}^{N} r_{g\,i}$$

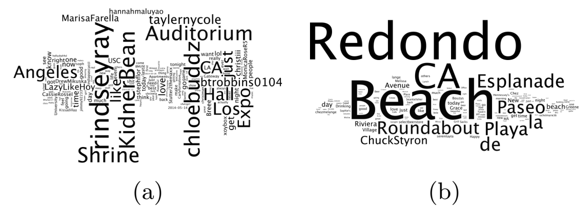

Figure 3: Word cloud of tweets from a single tract from (a) All LA Tracts and from (b) Tracts with >3 Check-ins.

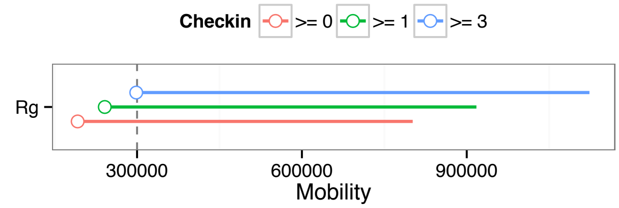

Figure 4: Distribution of user mobility by different groups of tracts. We measure mobility in a tract as average radius of gyration (in meters) of users tweeting from that tract.

where $r_{g\,i}$ is the radius of gyration of the ith of N users who tweet from that tract. This allows us to compare the human mobility of tracts with check-ins.

The Figure 4 plots the means and standard deviations (right tail) of R_g for the three groups of tracts containing check-ins, and the Table 1 complement these informations with the significance test between them. Based on the results, we reject the null hypothesis and conclude that the means of these distributions are significantly different from each other ($p < 0.001$ for all pairwise comparisons). We conclude that people tweeting from tracts with more check-ins travel farther, on average, than other people in our data set. Two distinct mechanisms could explain this difference. First, people may need to travel longer distances to go to the places with more popular amenities. This would suggest that amenities that people want to use are not located equitably, forcing people to travel greater distances to use them. Alternately, however, residents of tracts with more popular amenities could have higher mobility in general (perhaps, they can better afford to travel).

Demographics

We further examine whether differences between the three groups could be explained by demographics of their residents. For this analysis, we used three demographic variables widely related in the literature [3, 8, 10, 23] and available in the US Census per tract: age, ethnicity, and education. These variables are then linked to our results in human mobility and sentiment analysis.

We find that age by groups of tracts containing check-ins vary from 36 ($check-ins \geqslant 0$) to 38 years old ($check-ins \geqslant 3$). Ages are normally distributed and their differences seems small, but statistical inferences to test whether means of these distributions are the similar reveal the opposite ($p < 0.001$). In other words, these results show that residents

Table 1: Sentiment scores and mobility distance means organized by groups of check-ins. The first column represents the groups of tracts based on empirical check-in thresholds, and the respective quantity of tracts for each group in described in the second column. The subsequent columns relate the means, standard deviations and significance levels of Valence, Arousal and Dominance (WBK's lexicon), as well as of Positive and Negative (SentiStrength) sentiment scores for the given set of observed tracts. The last column presents the mean mobility distance (R_g) for each group.

		WKB Lexicon			**SentiStrength**		R_g
Check-ins	**# Tracts**	**Valence**	**Arousal**	**Dominance**	**Positive**	**Negative**	—
$\geqslant 0$	1718	4.776(0.361)***	3.493(0.236)***	4.632(0.335)***	1.587(0.089)**	1.410(0.075)***	191,9(311,1)***
$\geqslant 1$	687	4.869(0.351)***	3.534(0.224)***	4.712(0.321)***	1.597(0.080)**	1.392(0.073)***	241,2(345,0)***
$\geqslant 3$	162	4.989(0.292)***	3.589(0.179)**	4.814(0.259)***	1.603(0.072)ns	1.367(0.064)***	298,5(420,8)***

not significant (ns), P>0.10, $^+$P<0.10, *P<0.05, **P<0.01, ***P<0.001

in locations with more check-ins tend to be slightly older indeed.

Los Angeles attracts many Hispanic groups due to its strategic position and border, which makes this city an interesting case study. For this reason, we decided to test two different ethnicities available at the US Census: Hispanic and Non-Hispanic tracts. Inferential tests show at a high significance level ($p < 0.001$) that tracts with more check-ins have high Non-Hispanic population and low Hispanic population, suggesting that attractive amenities are located in places where fewer Hispanics live.

The last demographic measure we use is education, given by percentage of residents in a tract who have received bachelor's degree from a college or university, a master's, professional, or doctorate degrees. For each group, the percentage of educated people is : $check-ins \geqslant 0$ 19%, $check-ins \geqslant 1$ 24%, and $check-ins \geqslant 3$ 30%. The inference testing results show at a high significant level ($p < 0.001$) that differences in education exist between the three groups. Then, places containing geo-tagged tweets with more check-ins have indeed more educated residents.

6. DISCUSSIONS AND FUTURE WORKS

Using previous methods [18, 23, 27, 28, 34], in this paper we combined social media data from Twitter and Foursquare, with demographic data from the US Census to carry out a deeper micro-level analysis of geography and emotion within a city. Our results reveal that tracts with more check-ins are happier and have less negative geo-tagged tweets, which is fundamentally different from other areas within the Los Angeles County. These places offer amenities that people like to use, such as restaurants, parks, beaches, and gyms. The population of tracts with more check-ins tends to be slightly older, better educated, more Non-Hispanic, and more mobile, traveling farther, on average, than other people in our data set. This suggests that areas offering desirable amenities encourage people to commute longer distances to use them, although this observation may also be explained if residents of tracts with check-ins travelled more than other people.

Researchers have urged caution when using social media data, in particular Twitter, to study social science questions [33]. Twitter users may not be representative of the population researchers intended to study. Although we cannot eliminate all criticisms, we believe that our approach mitigates at least some of these concerns.

One fundamental limitation in our work is that we are not able to distinguish between residents and visitors to the area. While demographic analysis applies to residents, we have extended it to all people tweeting from the tract. Another concern is that Foursquare users are different, perhaps they are younger and better educated, so using them to select "attractive" tracts may skew the data. Further work is required to address these questions. However, even with these caveats, social media offers an intriguing data source for monitoring happiness in urban areas and exploring the questions of how happiness is connected to land use, which can in turn be used for city planners to design happier and more equitable cities

7. ACKNOWLEDGMENTS

This paper is partially funded by the CNPq (201224/2014–3), Brazil and by the Swiss National Science Foundation (CR21I1_146499). These supports are gratefully acknowledged by the authors.

References

[1] Happiness, geography and the environment. *Ecological Economics*, 65(2):386–396, Apr. 2008.

[2] A. Abbasi, A. Hassan, and M. Dhar. Benchmarking twitter sentiment analysis tools. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, 2014.

[3] A. Alshamsi, E. Awad, M. Almehrezi, V. Babushkin, P.-J. Chang, Z. Shoroye, A.-P. Toth, and I. Rahwan. Misery loves company: happiness and communication in the city. *EPJ Data Science*, 4(1):7, 2015.

[4] R. Alvarez, D. Garcia, Y. Moreno, and F. Schweitzer. Sentiment cascades in the 15m movement. *EPJ Data Science*, 4(1):1–13, 2015.

[5] Z. Cheng, J. Caverlee, K. Lee, and D. Sui. Exploring millions of footprints in location sharing services. In *International AAAI Conference on Web and Social Media*, 2011.

[6] H. Cramer, M. Rost, and L. E. Holmquist. Performing a check-in: Emerging practices, norms and 'conflicts' in location-sharing using foursquare. In *Proceedings of the 13th International Conference on Human Computer Interaction with Mobile Devices and Services*, 2011.

[7] S. Duffy. A community creating their own rules on foursquare. In *Proceedings of the First NSS Workshop on Mobile and Online Social Networks*, 2011.

[8] N. Eagle, M. Macy, and R. Claxton. Network diversity and economic development. *Science*, 328(5981):1029–1031, 2010.

[9] N. Eagle and A. Pentland. Reality mining: Sensing complex social systems. *Personal Ubiquitous Computing*, 10(4):255–268, Mar. 2006.

[10] R. A. Easterlin. Does economic growth improve the human lot? some empirical evidence. *Nations and households in economic growth*, 89:89–125, 1974.

[11] J. R. Fontaine, K. R. Scherer, E. B. Roesch, and P. C. Ellsworth. The world of emotions is not two-dimensional. *Psychological science*, 18(12):1050–1057, 2007.

[12] A. Garas, D. Garcia, M. Skowron, and F. Schweitzer. Emotional persistence in online chatting communities. *Scientific Reports*, 2, 2012.

[13] D. Garcia, A. Garas, and F. Schweitzer. Positive words carry less information than negative words. *EPJ Data Science*, 1(1):3, 2012.

[14] D. Garcia, F. Mendez, U. Serdült, and F. Schweitzer. Political polarization and popularity in online participatory media: an integrated approach. In *Proceedings of the first edition workshop on Politics, elections and data*, 2012.

[15] D. Garcia and F. Schweitzer. Emotions in product reviews–empirics and models. In *Privacy, Security, Risk and Trust (PASSAT) and 2011 IEEE Third Inernational Conference on Social Computing (SocialCom), 2011 IEEE Third International Conference on*, 2011.

[16] S. A. Golder and M. W. Macy. Diurnal and seasonal mood vary with work, sleep, and daylength across diverse cultures. *Science*, 333(6051):1878–1881, Sept. 2011.

[17] P. Gonçalves, M. Araújo, F. Benevenuto, and M. Cha. Comparing and combining sentiment analysis methods. In *Proceedings of the first ACM conference on Online social networks*, 2013.

[18] M. C. González, C. A. Hidalgo, and A.-L. Barabási. Understanding individual human mobility patterns. *Nature*, 453(7196):779–782, June 2008.

[19] S. González-Bailón, R. E. Banchs, and A. Kaltenbrunner. Emotions, public opinion, and us presidential approval rates: A 5-year analysis of online political discussions. *Human Communication Research*, 38(2):121–143, 2012.

[20] F. Kivran-Swaine and M. Naaman. Network properties and social sharing of emotions in social awareness streams. In *Proceedings of the ACM 2011 Conference on Computer Supported Cooperative Work*, 2011.

[21] A. Kramer. An unobtrusive behavioral model of "gross national happiness". In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2010.

[22] O. Kucuktunc, B. B. Cambazoglu, I. Weber, and H. Ferhatosmanoglu. A large-scale sentiment analysis for yahoo! answers. In *Proceedings of the fifth ACM international conference on Web search and data mining*, 2012.

[23] L. Mitchell, M. R. Frank, K. D. Harris, P. S. Dodds, and C. M. Danforth. The geography of happiness: Connecting twitter sentiment and expression, demographics, and objective characteristics of place. *PLoS ONE*, 8(5):e64417, 05 2013.

[24] A. Noulas, S. Scellato, C. Mascolo, and M. Pontil. An empirical study of geographic user activity patterns in foursquare. In *Proceedings of the 5th Int'l AAAI Conference on Weblogs and Social Media*, 2011.

[25] C. E. Osgood, G. J. Suci, and P. H. Tannenbaum. *The measurement of meaning.* University of Illinois Press, 1964.

[26] B. Pang and L. Lee. Opinion mining and sentiment analysis. *Foundations and trends in information retrieval*, 2(1-2):1–135, 2008.

[27] D. Quercia, J. Ellis, L. Capra, and J. Crowcroft. Tracking "gross community happiness" from tweets. In *Proceedings of the CSCW'12 Conference*, 2012.

[28] D. Quercia and D. Saez-Trumper. Mining urban deprivation from foursquare: Implicit crowdsourcing of city land use. *Pervasive Computing, IEEE*, 13(2):30–36, Apr 2014.

[29] A. Rajadesingan, R. Zafarani, and H. Liu. Sarcasm detection on twitter: A behavioral modeling approach. In *Proceedings of the Eighth ACM International Conference on Web Search and Data Mining*, 2015.

[30] J. A. Russell and A. Mehrabian. Evidence for a three-factor theory of emotions. *Journal of research in Personality*, 11(3):273–294, 1977.

[31] M. Thelwall, K. Buckley, and G. Paltoglou. Sentiment strength detection for the social web. *Journal of the American Society for Information Science and Technology*, 63(1):163–173, 2012.

[32] M. Thelwall, K. Buckley, G. Paltoglou, M. Skowron, D. Garcia, S. Gobron, J. Ahn, A. Kappas, D. Küster, and J. A. Holyst. Damping sentiment analysis in online communication: discussions, monologs and dialogs. In *Computational linguistics and intelligent text processing*, pages 1–12. Springer, 2013.

[33] Z. Tufekci. Big questions for social media big data: Representativeness, validity and other methodological pitfalls. 2014.

[34] A. Venerandi, G. Quattrone, D. Quercia, L. Capra, and D. Saez-Trumper. Measuring urban deprivation from user generated content. In *Proceedings of the CSCW'15 Conference*, 2015.

[35] A. B. Warriner, V. Kuperman, and M. Brysbaert. Norms of valence, arousal, and dominance for 13,915 english lemmas. *Behavior research methods*, 45(4):1191–1207, 2013.

[36] D. Watson, L. A. Clark, and A. Tellegen. Development and validation of brief measures of positive and negative affect: the panas scales. *Journal of personality and social psychology*, 54(6):1063, 1988.