

Emotions, Demographics and Sociability in Twitter Interactions

Kristina Lerman

USC Information Sciences Institute
Marina del Rey, CA 90292
lerman@isi.edu

Megha Arora

IIT-Delhi
New Delhi, India - 110020
megha12059@iiitd.ac.in

Luciano Gallegos

USC Information Sciences Institute
Marina del Rey, CA 90292
luciano.gallegos@gmail.com

Ponnurangam Kumaraguru

IIT-Delhi
New Delhi, India - 110020
pk@iiitd.ac.in

David Garcia

ETH-Zurich
Zurich, Switzerland
dgarcia@ethz.ch

Abstract

The social connections people form online affect the quality of information they receive and their online experience. Although a host of socioeconomic and cognitive factors were implicated in the formation of offline social ties, few of them have been empirically validated, particularly in an online setting. In this study, we analyze a large corpus of geo-referenced messages, or tweets, posted by social media users from a major US metropolitan area. We linked these tweets to US Census data through their locations. This allowed us to measure emotions expressed in the tweets posted from an area, the structure of social connections, and also use that area's socioeconomic characteristics in analysis. We find that at an aggregate level, places where social media users engage more deeply with less diverse social contacts are those where they express more negative emotions, like sadness and anger. Demographics also has an impact: these places have residents with lower household income and education levels. Conversely, places where people engage less frequently but with diverse contacts have happier, more positive messages posted from them and also have better educated, younger, more affluent residents. Results suggest that cognitive factors and offline characteristics affect the quality of online interactions. Our work highlights the value of linking social media data to traditional data sources, such as US Census, to drive novel analysis of online behavior.

Introduction

Humans have evolved large brains, in part to handle the cognitive demands of social relationships (Dunbar 2003). The social structures resulting from these relationships confer numerous fitness advantages. Scholars distinguish between two types of social relationships: those representing strong and weak ties. Strong ties are characterized by high frequency of interaction and emotional intimacy that can be found in relationships between family members or close friends. People connected by strong ties share mutual friends (Granovetter 1973), forming cohesive social bonds that are essential for providing emotional and material support (Putnam 2000; Rimé 2009) and creating resilient communities (Sampson, Raudenbush, and Earls 1997). In contrast, weak ties represent more casual social relationships, characterized by less frequent, less intense interactions, such

as those occurring between acquaintances. By bridging otherwise unconnected communities, weak ties expose individuals to novel and diverse information that leads to new job prospects (Granovetter 1983) and career opportunities (Burt 1995; Burt 2004). Online social relationships provide similar benefits to those of the offline relationships, including emotional support and exposure to novel and diverse information (Aral and Van Alstyne 2011; Bakshy et al. 2012; Kang and Lerman 2015).

How and why do people form different social ties, whether online or offline? Of the few studies that addressed this question, Shea et al. (Shea et al. 2015) examined the relationship between emotions and cognitive social structures, i.e., the mental representations individuals form of their social contacts (Krackhardt 1987). In a laboratory study, they demonstrated that subjects experiencing positive affect, e.g., emotions such as happiness, were able to recall a larger number of more diverse and sparsely connected social contacts than those experiencing negative affect, e.g., sadness. In other words, they found that positive affect was more closely associated with weak ties and negative affect with strong ties in cognitive social structures. This is consistent with findings that negative emotional experiences are shared more frequently through strong ties (Rimé 2009), not only to seek support but also as a means of strengthening the tie (Niedenthal and Brauer 2012). In addition to psychological factors, social structures also depend on the participants' socioeconomic and demographic characteristics. A study, which reconstructed a national-scale social network from the phone records of people living in the United Kingdom, found that people living in more prosperous regions formed more diverse social networks, linking them to others living in distinct communities (Eagle, Macy, and Claxton 2010). On the other hand, people living in less prosperous communities formed less diverse, more cohesive social structures.

The present paper examines how psychological and demographic factors affect the structure of online social interactions. We restrict our attention to interactions on the Twitter microblogging platform. To study these interactions, we collected a large body of geo-referenced text messages, known as tweets, from a large US metropolitan area. Further, we linked these tweets to US Census tracts through their locations. Census *tracts* are small regions, on a scale of city

blocks, that are relatively homogeneous with respect to population characteristics, economic status, and living conditions. Some of the tweets also contained explicit references to other users through the ‘@’ mention convention, which has been widely adopted on Twitter for conversations. We used mentions to measure the strength of social ties of people tweeting from each tract. Using these data we studied (at tract level) the relationship between social ties, the socioeconomic characteristics of the tract, and the emotions expressed by people tweeting from that tract. In addition, people tweeting from one tract often tweeted from other tracts. Since geography is a strong organizing principle, for both offline (Travers and Milgram 1969; Barthlemy 2011) and online (Quercia et al. 2012; Liben-Nowell et al. 2005; Backstrom, Sun, and Marlow 2010) social relationships, we measured the spatial diversity of social relationships, and studied its dependence on socioeconomic, demographic, and psychological factors.

Our work complements previous studies of offline social networks and demonstrates a connection between the structure of online interactions in urban places and their socioeconomic characteristics. More importantly, it links the structure of online interactions to positive affect. People who express happier emotions interact with a more diverse set social contacts, which puts them in a position to access, and potentially take advantage of, novel information. As our social interactions increasingly move online, understanding, and being able to unobtrusively monitor, online social structures at a macroscopic level is important to ensuring equal access to the benefits of social relationships.

In the rest of the paper, we first describe data collection and methods used to measure emotion and social structure. Then, we present results of a statistical study of social ties and their relationships to emotions and demographic factors. The related works are addressed after this. Although many important caveats exist about generalizing results of the study, especially to offline social interactions, our work highlights the value of linking social media data to traditional data sources, such as US Census, to drive novel analysis of online behavior and online social structures.

Related Work

Eagle et al. (Eagle, Macy, and Claxton 2010) explored the link between socioeconomic factors and network structure using anonymized phone call records to reconstruct the national-level network of people living in the UK. Measures of socioeconomic development were constructed from the UK government’s Index of Multiple Deprivation (IMD), a composite measure of prosperity based on income, employment, education, health, crime, housing of different regions within the country. They found that people living in more prosperous regions formed more diverse social networks, linking them to others living in distinct communities. On the other hand, people living in less prosperous communities formed less diverse, more cohesive social structures.

Quercia et al. (Quercia et al. 2012) found that sentiment expressed in tweets posted around 78 census areas of London correlated highly with community socioeconomic well being, as measured by the Index of Multiple Deprivation

(i.e., qualitative study of deprived areas in the UK local councils). In another study (Quercia, Capra, and Crowcroft 2012) they found that happy places tend to interact with other happy places, although other indicators such as demographic data and human mobility were not used in their research (Cheng et al. 2011).

Other researcher used demographic factors and associated them to sentiment analysis to measure happiness in different places. For instance, Mitchell et al. (Mitchell et al. 2013) 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.

Psychological and cognitive states affect the types of social connections people form and their ability to recall them (Brashears 2013). When people experience positive emotions, or affect, they broaden their cognitive scope, widening the array of thoughts and actions that come to mind (Fredrickson 2001). In contrast, experiencing negative emotions narrow attention to the basic actions necessary for survival. Shea et al. (Shea et al. 2015) tested these theories in a laboratory, examining the relationship between emotions and the structure of networks people were able to recall. They found that subjects experiencing positive affect were able to recall a larger number of more diverse and sparsely connected social contacts than those experiencing negative emotions. The study did not resolve the question of how many of the contacts people were able to recall that they proceeded to actively engage.

A number of innovative research works attempted to better understand human emotion and mobility. Some of these works focuses on geo-tagged location data extracted from Foursquare and Twitter. Researchers reported (Cramer, Rost, and Holmquist 2011; Noulas et al. 2011) 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 (Duffy 2011).

Methods

Data

Los Angeles (LA) County is the most populous county in the United States, with almost 10 million residents. It is extremely diverse both demographically and economically, making it an attractive subject for research. We collected a large body of tweets from LA County over the course of 4 months, starting in July 2014. Our data collection strategy was as follows. 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 (time-

line) tweets from users who tweeted from within this area during this time period. A portion of these tweets were geo-referenced, i.e. they had geographic coordinates attached to them. In all, we collected 6M geo-tagged tweets made by 340K distinct users.

We localized geo-tagged tweets to tracts from the 2012 US Census.¹ A tract is a geographic region that is 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 demographic characteristics of that population. We included only Los Angeles County tracts in the analysis. We used data from the US Census to obtain demographic and socioeconomic characteristics of a tract, including the mean household income, median age of residents, percentage of residents with a bachelor’s degree or above, as well as racial and ethnic composition of the tract.

Emotion Analysis

To measure emotions, we apply sentiment analysis (Pang and Lee 2008), i.e. methods that process text to quantify subjective states of the author of the text. Two recent independent benchmark studies evaluate a wide variety of sentiment analysis tools in various social media (Gonçalves et al. 2013) and Twitter datasets (Abbasi, Hassan, and Dhar 2014). Across social media, one of the best performing tools is SentiStrength (Thelwall, Buckley, and Paltoglou 2012), which also was shown to be the best unsupervised tool for tweets in various contexts (Abbasi, Hassan, and Dhar 2014).

SentiStrength quantifies emotions expressed in short informal text by matching terms from a lexicon and applying intensifiers, negations, misspellings, idioms, and emoticons. We use the standard English version of SentiStrength² to each tweet in our dataset, quantifying positive sentiment $P \in [+1, +5]$ and negative sentiment $N \in [-1, -5]$, consistently with the Positive and Negative Affect Schedule (PANAS) (Watson, Clark, and Tellegen 2013). SentiStrength has been shown to perform very closely to human raters in validity tests (Thelwall, Buckley, and Paltoglou 2012) and has been applied to measure emotions in product reviews (Garcia and Schweitzer 2011), online chatrooms (Garas et al. 2012), Yahoo answers (Kucuktunc et al. 2012), and Youtube comments (Garcia et al. 2012). In addition, SentiStrength allows our approach to be applied in the future to other languages, like Spanish (Alvarez et al. 2015), and to include contextual factors (Thelwall et al. 2013), like sarcasm (Rajadesingan, Zafarani, and Liu 2015).

Beyond positivity and negativity, meanings expressed through text can be captured through the application of the semantic differential (Osgood, Suci, and Tannenbaum 1964), a dimensional approach that quantifies emotional meaning in terms of valence, arousal, and dominance (Russell and Mehrabian 1977). The dimension of *valence* quantifies the level of pleasure or evaluation 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 or potency experienced in relation to an emotional word. Research in psychology suggests that a multidimensional approach is necessary to capture the variance of emotional experience (Fontaine et al. 2007), motivating our three-dimensional measurement beyond simple polarity approximations. The state of the art in the quantification of these three dimensions is the lexicon of Warriner, Kuperman, and Brysbaert (WKB) (Warriner, Kuperman, and Brysbaert 2013). The WKB lexicon includes scores in the three dimensions for more than 13,000 English lemmas. We quantify these three dimensions in a tweet by first lemmatizing the words in the tweet, to then match the lexicon and compute mean values of the three dimensions as in (González-Bailón, Banchs, and Kaltenbrunner 2012). The large size of this lexicon allows us to match terms in 82.39% of the tweets in our dataset, which we aggregate to produce multidimensional measures of emotions.

The Figure 1 presents word clouds of tweets from a tract with one of the highest average valence and one from a tract with a lower average valence. The words themselves are colored by their valence, with red corresponding to high and blue to low valence words. Despite seemingly small differences in average tract valence, the words depicted in the word clouds are remarkably different in the emotions they convey. The “happy” tract has words such as ‘beach’, ‘love’, ‘family’, ‘beautiful’, while the “sad” tract contains many profanities (though it also contains some happy words).

Social Tie Analysis

Twitter users address others using the ‘@’ mention convention. We use the mentions as evidence of social ties, although sometimes users address public figures and celebrities also using this convention. We use mention frequency along a tie as a proxy of tie strength, drawing upon multiple studies that used frequency of interactions as a measure of tie strength (Granovetter 1983; Onnela et al. 2007; Quercia et al. 2012). In contrast to other measures, such as clustering coefficient, it does not require knowledge of full network structure (which we do not observe).

Tie strength For each tract, we create a mention graph with users as nodes and an edge from user A to user B if A mentions B in her tweets. Using this graph, the average social tie strength per tract is defined as

$$S_i = \frac{\sum_{j=1}^{k_i} w_j}{k_i} \quad (1)$$

where w_j is the weight of the j^{th} edge (i.e., the number of times user A mentioned user B), and k_i is the total number of distinct users mentioned in tract i .

We do not have complete knowledge of network structure, since we only observe the tweets of users who geo-referenced their tweets, and not necessarily the tweets of mentioned users. However, even in the absence of complete information about interactions, average tie strength captures the amount of social cohesion and diversity (Gonçalves, Perra, and Vespignani 2011). Figure 2 illustrates mention

¹American Fact Finder (<http://factfinder.census.gov/>)

²<http://sentistrength.wlv.ac.uk/>

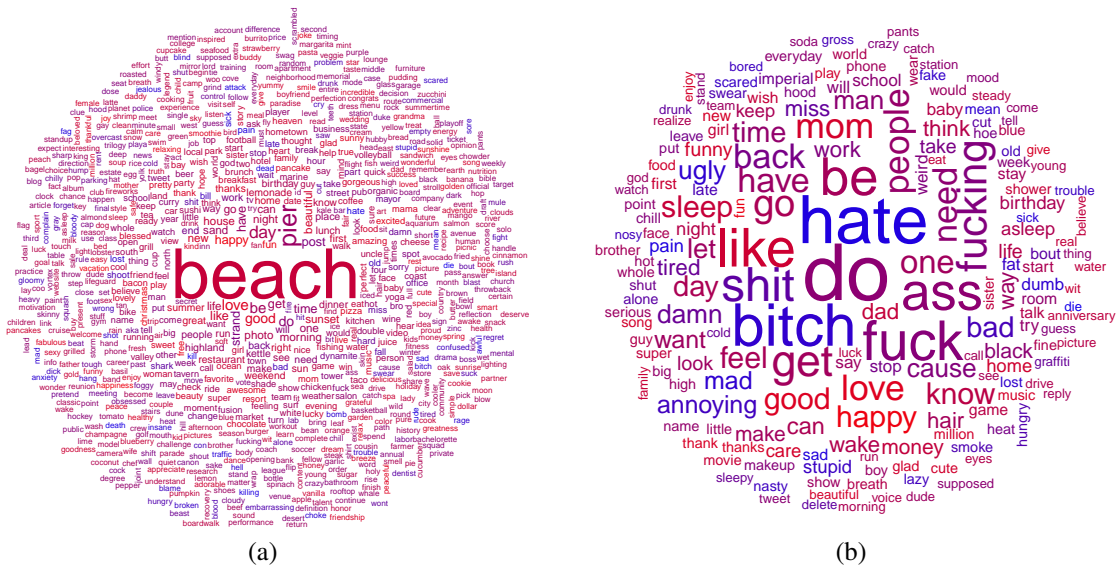


Figure 1: Word cloud of tweets from two tracts with average valence of (a) 6.122 and (b) 5.418. Words are colored by their valence, with red corresponding to high valence words, and blue to low valence words.

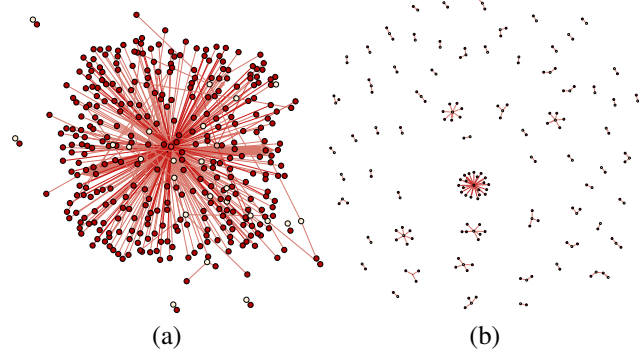


Figure 2: Mentions graphs of two different tracts showing (a) strong ties (average tie strength $S_i = 7.33$) and (b) weak ties ($S_i = 1.08$). Tweeting users are represented as white nodes, while mentioned users are pink in color. The width of an edge represents the number of mentions.

graphs from two tracts with very different tie strength values. High tie strength (Fig. 2(a)) is associated with a high degree of interaction and more clustering (Granovetter 1973). In contrast, low tie strength is associated with a sparse, more diverse network with few interconnections (Fig 2(b)).

Spatial diversity Geography and distance are important organizing principles of social interactions, both offline (Travers and Milgram 1969; Barthlemy 2011) and online (Quercia et al. 2012; Liben-Nowell et al. 2005; Backstrom, Sun, and Marlow 2010). While most social interactions are short-range, long-distance interactions serve as evidence of social diversity (Eagle, Macy, and Claxton 2010).

In this paper, we use the movement of people across tracts as evidence of the spatial diversity of their social structures. Following Eagle et al. (Eagle, Macy, and Claxton 2010), we measure spatial diversity of places from which people tweeting from a given tract also tweet from, using Shannon’s Entropy ratio, as

$$D_i = \frac{-\sum_{j=1}^{n_i} p_{ij} \log(p_{ij})}{\log n_i} \quad (2)$$

where n_i is the number of tracts from which users who tweeted from tract i also tweeted from, and p_{ij} is the proportion of tweets posted by these users from tract j such that

$$p_{ij} = \frac{T_{ij}}{\sum_{j=1}^{n_i} T_{ij}} \quad (3)$$

where T_{ij} is the number of tweets that have been posted in tract j by the users who have tweeted from both tract i and j .

Thus, spatial diversity is a ratio that compares the empirical entropy of data with its expected value in the uniformly distributed case. As a consequence, a high spatial diversity value for a tract suggests that people tweeting from that tract split their tweets evenly among all the tracts they are tweeting from. In contrast, a low value implies that people tweeting from that tract concentrate their tweets in few tracts.

Results

Happier places tend to attract more Twitter users: the correlation between the mean valence of tweets posted from a tract and the number of people tweeting from the tract is 0.30 ($p < 0.001$). Beyond this correlation, we observe systematic trends between emotions expressed in tweets posted from a tract, and the structure of social interactions of people tweeting from that tract.

A number of variables seem to contribute to shaping the structure of online social interactions within an urban area. Amongst these possibilities are *affect* and *demographic* factors. The *affect* variables are: valence (V), arousal (A), dominance (D), positive (P), negative (N) sentiment. The *demographic* variables are: mean income (inc) of population within a tract, education (edu) as measured by the percentage of population with bachelor’s degree or above, percentage of employed residents (emp), their median age (age), and the fraction of Hispanic (hsp), Asian (asi), African American (blk), populations white (wht) populations within a tract. The correlations of these variables with tie strength (tie) and spatial diversity (spa) are described in Tables 1 and 2.

Although not presented in these Tables, another interesting relation is between valence and dominance: they are highly and significantly correlated ($r = 0.88, p < 0.001$). Bearing these relations in mind, we test linear regression models for tie strength ($tie \sim V \times D + A, tie \sim P + N, tie \sim inc + edu + age + emp$ and $tie \sim hsp + wht + blk + asi$) and spatial diversity ($spa \sim V \times D + A, spa \sim P + N, spa \sim inc + edu + age + emp$ and $spa \sim hsp + wht + blk + asi$).

Regression coefficients for each model are computed on normalized data and summarized in Figure 3. For tie strength (top row), some interesting results are depicted. In (a), while arousal and valence:dominance are mainly neutral, dominance and valence have negative dependencies with strength of ties: i.e., the lower the valence, the stronger the tie strength. Similarly, in (b), negative and positive sentiment represent negative dependencies with tie strength. Hence, larger positive sentiment (P) values are associated with weaker social ties. Stronger ties are associated with smaller values of N , i.e., more negative sentiment. In (c), education has a negative dependence, while others are basically neutral and, in (d), Hispanic ethnicity has a positive dependence to tie strength, with others basically neutral.

Regarding spatial diversity (bottom row in Fig. 3), results are more diverse. In (e), dominance and valence:dominance are mainly neutral, whereas arousal presents negative dependence and valence presents positive dependence, i.e. the higher the valence, the broader the spatial diversity. In (f), negative and positive sentiment have positive dependencies with respect to spatial diversity. In (g), age is neutral, education and employment percentage present positive and mean income negative dependencies. Finally, in (h), besides Hispanic presenting negative dependence, all others are mainly neutral.

Overall, these results show that besides the correlations between the chosen variables, some of them contribute more to explain the tie strength and social diversity while others have minor contributions. More details regarding affect, demographics and social interactions results are presented in the following.

Emotion and Social Ties

Table 1 reports correlations between affect and the structure of social interactions measured from tweets. We quantify affect using mean valence, arousal and dominance of tweets posted from a tract (measured by WKB lexicon),

and mean positive and negative sentiment, (measured by SentiStrength). The social variables are average social tie strength per tract and spatial diversity, a measure of inter-tract mobility. Most of these correlations are statistically significant.

Variables	Tie Strength	Spatial Diversity
Valence	-0.36***	0.32***
Arousal	0.14***	-0.20***
Dominance	-0.31***	0.31***
Positive Sent.	-0.18***	0.23***
Negative Sent.	-0.24***	0.28***

* $p_i 0.05$, ** $p_i 0.01$, *** $p_i 0.001$

Table 1: Correlation coefficient of per-tract strength of ties measure and mean value of affect.

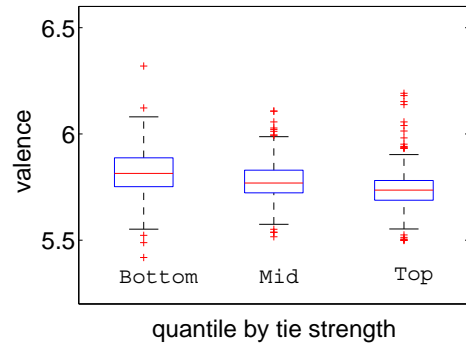


Figure 4: Box plot of mean valence of tweets posted from tracts, after they were grouped by mean tie strength into three equal-sized groups, or tertiles. Lines inside the boxes indicate median value of valence, whereas the boxes themselves show interquartile ranges.

The average value of valence across all tracts is 5.78, which corresponds to slightly positive affect with respect to the neutral point of 5.0, in line with emotional expression in other media (Garcia, Garas, and Schweitzer 2012). This contrasts with the negative correlation between valence and tie strength. Consequently, tracts with stronger social ties are associated with less positive — sadder, angrier — tweets than tracts with weaker social ties, which are associated with more positive — happier — tweets. This can be seen clearly in Figure 4, which shows the values of mean valence for tracts after they were ordered by average social tie strength and divided into three equal-sized bins, or tertiles. The bottom third of tracts, i.e., those with weakest social ties, have the highest values of valence (mean 5.82). In contrast, the top tertile composed of tracts with strongest social ties has the lowest values of valence (mean 5.74). These differences are statistically significant ($p < 0.01$).

For a different perspective on this relationship, we order tracts by the average valence, arousal, and dominance of

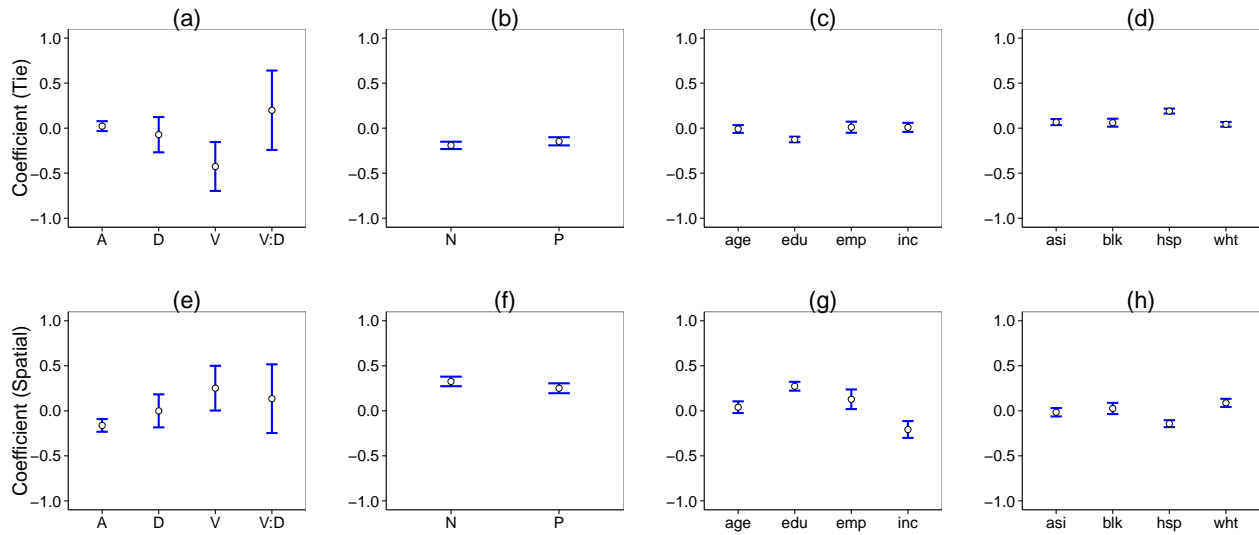


Figure 3: Error bars based on linear regression models: coefficients (tie, spa) are estimated against sentiment (WBK (a), SentiStrength (b)), socioeconomics (c) and ethnicity (d) variables

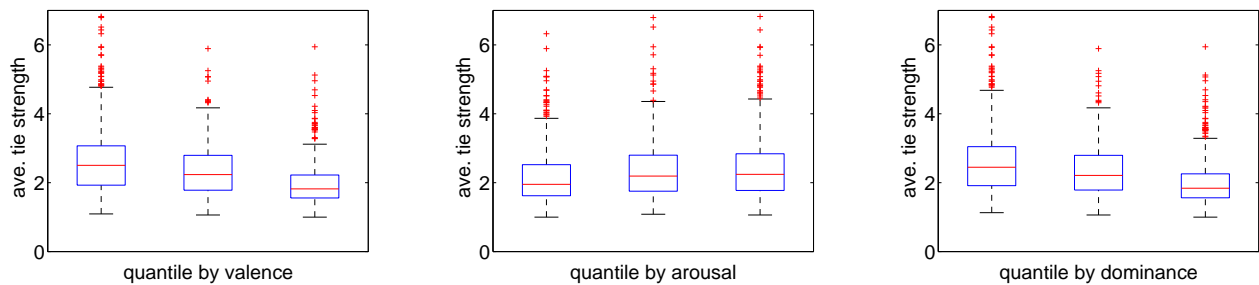


Figure 5: Social tie strength analysis. Tracts were ordered by the mean value of (a) valence, (b) arousal, or (c) dominance of tweets from that tract and divided into three groups of equal size. Lines inside the boxes indicate median value of tie strength for each group, whereas the boxes show interquartile ranges.

their tweets and group them into three equal-size bins. Figure 5 reports the average tract tie strength across these tertiles. Tracts from which people post lower valence (sadder) tweets, on average, are associated with significantly stronger social ties ($p < 0.01$) than tracts with higher valence (happier) tweets (Figure 5(a)). Tracts from which users post messages expressing higher arousal tend to have slightly stronger ties (Figure 5(b)) although only the difference between mean tie strength of the mid and top tertiles are significant ($p < 0.05$). In contrast, tracts associated with message expressing higher dominance, Figure 5(b), are associated with significantly weaker social ties ($p < 0.01$). Since weak ties are associated with more diverse social relationships, our observation confirms the relationship between emotions and diversity of network structure.

The sentiment values computed by SentiStrength are consistent with this trend. The correlation between tie strength and positive sentiment while weaker, is also negative, suggesting that tracts with more positive tweets have weaker

ties. Note that values for negative sentiment are below zero, with lower values representing stronger negative sentiment. In this case, negative correlation with tie strength means that tracts with more negative tweets have stronger ties than tracts with less negative tweets. This result is in line with theories of social regulation of emotions (Rimé 2009) and with previous results in protest movements that showed how online negative emotions were associated with stronger collective action (Alvarez et al. 2015).

We also report the relationship between affect of a tract and the spatial diversity of users tweeting from that tract, a measure of inter-tract mobility defined in Section . As shown in Figure 6, tracts with more positive tweets have significantly higher spatial diversity than tracts with less positive tweets ($p < 0.001$). This suggests that people expressing happier emotions move (and tweet from) a larger number of different places than people expressing more negative emotions, whose movements are confined to a smaller set of tracts.

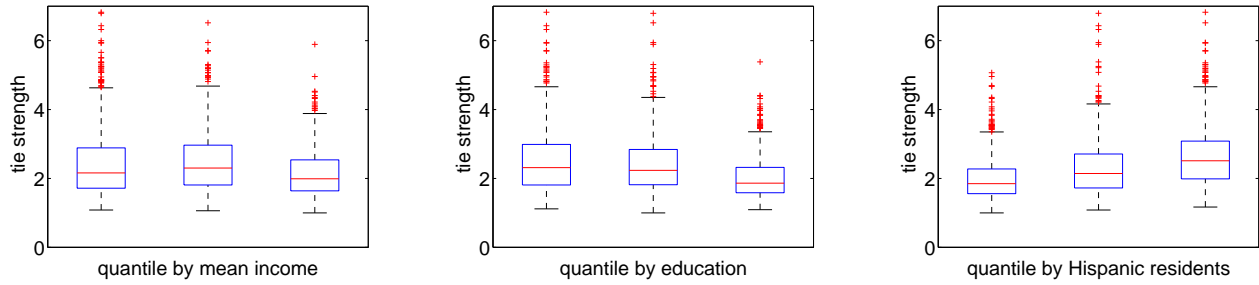


Figure 7: Demographic analysis of social ties. Tracts were ordered by demographic characteristic (mean household income, percentage of residents with Bachelor’s degree or above, and number of Hispanic residents) in that tract and divided into three groups of equal size. Lines inside the boxes indicate median value of tie strength within each group, whereas the boxes show interquartile ranges. Leftmost box represents the bottom third of tracts with lowest values of the demographic attribute, and rightmost box represents the top third of tracts with largest values of the attribute.

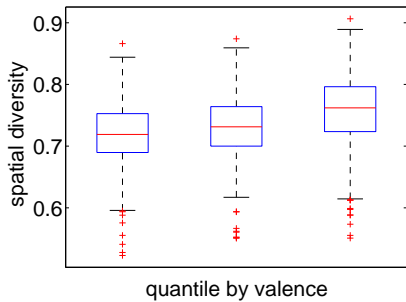


Figure 6: Spatial (inter-tract) diversity for three quantiles, ordered by the mean valence of tweets posted from them.

Demographics and Social Ties

We find a relationship between demographic characteristics of tracts from which people tweet and the social structure of tweeting users. Table 2 reports correlations between average per-tract social tie strength, spatial diversity and demographic variables extracted from US Census data. While a previous study (Eagle, Macy, and Claxton 2010) of the relationship between social structure and economic prosperity used somewhat different indicators — e.g., they used network diversity to measure diversity of social contacts — we reach qualitatively similar conclusions. Specifically, we find a negative correlation between (mean) income, one measure of economic prosperity, and sociability: higher incomes are associated with weaker social ties. There is an even stronger negative relationship between education (measured by the percentage of tract residents who hold a Bachelor’s degree or higher) and sociability. Socioeconomic variables are highly inter-correlated, and it is difficult to infer causal relationships within data. Still, the higher correlation of education with tie strength and mobility implies that education may be more predictive of sociability than other demographic attributes.

We take a closer look at the relationships between social tie strength and demographic attributes using quantile

Variables	Tie Strength	Spatial Diversity
Mean Income	-0.12***	0.08**
Employment	-0.12***	0.23***
Age	-0.18***	0.18***
Education	-0.27***	0.35***
White	-0.14***	0.20***
African Amer.	0.06*	0.03
Hispanic	0.35***	-0.27***
Asian Amer.	-0.03	0.10***

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 2: Correlation between demographic characteristics of tracts and measures of social structure of people tweeting from these tracts.

analysis. Figure 7 shows the box plot of tie strength of a group of tracts after they were divided into three equal-size groups based on each demographic attribute. For example, Figure 7(top) divides tracts by mean income into the lowest-earning, mid-earning, and highest-earning tertiles and shows the range of social tie strength values in each tertile. Only the top income and education tertiles have significantly weaker social ties than the other two tertiles ($p < 0.01$). However, when tracts are divided into tertiles according to the number of Hispanic residents, the differences between the mean tie strengths of all tertiles are significant ($p < 0.01$). This result highlights the social component of culture: Hispanic cultures focus more on sociability values and are less individualist than anglo-saxon cultures (for example, Mexico scores 30 and the US 91 in the individualism scale of Hofstede (Hofstede 1980)). This provides an explanation for the stronger links of tracts with higher number of Hispanic residents, as their online network structures reflect their shared values (Garcia-Gavilanes, Quercia, and Jaimes 2013; Kayes et al. 2015).

Figure 8 shows the association between spatial diversity and demographic characteristics. Income does not appear to significantly affect spatial diversity: only the top tertile of

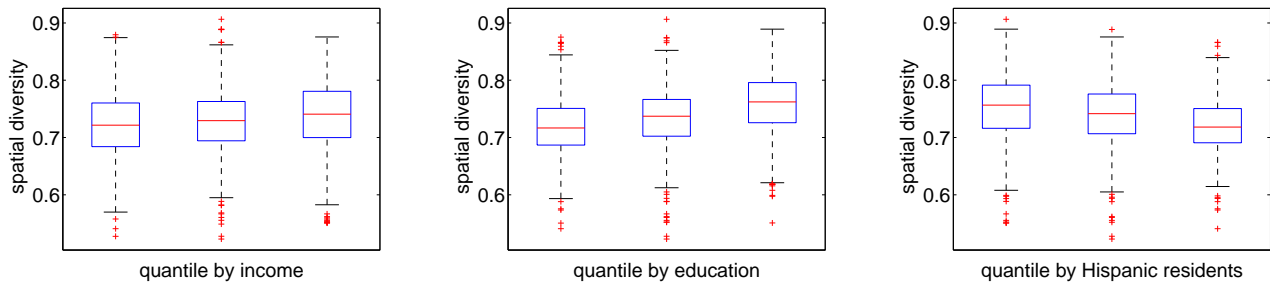


Figure 8: Demographic analysis of spatial diversity (inter-tract mobility). Tracts were ordered by (a) mean household income, (b) percentage of residents with Bachelor’s degree or above, or (c) number of Hispanic residents in that tract and divided into three groups of equal size. The figure shows the box plot of inter-tract mobility values of tracts in different quantiles. Leftmost box represents the bottom third of tracts with lower values of the demographic attribute, and rightmost box represents the top third of tracts with largest values of the attribute.

tracts by incomes has a significantly different spatial diversity ($p < 0.001$) from the other two tertiles. Education, however, has a stronger dependence: tracts with better-educated residents also have significantly higher ($p < 0.001$) spatial diversity than tracts with fewer educated residents. In addition, ethnicity appears to be a factor. Tracts with larger Hispanic population have significantly lower spatial diversity ($p < 0.01$) than other tracts.

Discussions

The availability of large scale, near real-time data from social media sites such as Twitter brings novel opportunities for studying online behavior and social interactions at an unprecedented spatial and temporal resolution. By combining Twitter data with US Census, we were able to study how the socioeconomic and demographic characteristics of residents of different census tracts are related to the structure of online interactions of users tweeting from these tracts. Moreover, sentiment analysis of tweets originating from a tract revealed a link between emotions and sociability of Twitter users.

Our findings are broadly consistent with results of previous studies carried out in an offline setting, and also give new insights into the structure of online social interactions. We find that at an aggregate level, areas with better educated, somewhat younger and higher-earning population are associated with weaker social ties and greater spatial diversity (or inter-tract mobility). In addition, Twitter users express happier, more positive emotions from these areas. Conversely, areas that have more Hispanic residents are associated with stronger social ties and lower spatial diversity. People also express less positive, sadder emotions in these areas. Since weak ties are believed to play an important role in delivering strategic, novel information, our work identifies a social inequity, wherein the already privileged ones (more affluent, better educated, happier) are in network positions that potentially allow them greater access to novel information.

Some important considerations limit the interpretation of our findings. First, our methodology for identifying social interactions may not give a complete view of the social net-

work of Twitter users. Our observations were limited to social interactions initiated by users who geo-reference their tweets. This may not be representative of all Twitter users posting messages from a given tract, if systematic biases exist in what type of people elect to geo-reference their tweets. For demographic analysis, we did not resolve the home location of Twitter users. Instead, we assumed that characteristics of an area, i.e., of residents of a tract, influence the tweets posted from that tract. Other subtle selection biases could have affected our data and the conclusions we drew (Tufekci 2014). It is conceivable that Twitter users residing in more affluent areas are less likely to use the geo-referencing feature, making our sample of Twitter users different from the population of LA county residents. Recognizing this limitation, we did not make any claims about the behavior of LA residents; rather, we focused on the associations between emotions and characteristics of a place and the behavior of Twitter users, with an important caveat that those who turn on geo-referencing may differ from the general population of Twitter users.

For the analysis of emotions, we only considered English language tweets, although a significant fraction of tweets were in Spanish. This may bias the average affect of tracts, especially for low-valence tracts, which have a larger number of Hispanic residents. In the future, we plan to address this question by conducting sentiment analysis of Spanish language tweets.

Acknowledgments

MA was supported by the USC Viterbi-India internship program. LG acknowledge support by the National Counsel of Technological and Scientific Development — CNPq, Brazil (201224/2014–3), and USC-ISI visiting researcher fellowship. This work was also partially supported by DARPA, under contract W911NF-12-1-0034. This support is gratefully acknowledged.

References

[Abbasi, Hassan, and Dhar 2014] Abbasi, A.; Hassan, A.; and Dhar, M. 2014. Benchmarking twitter sentiment analy-

- sis tools. In *LREC'14*).
- [Alvarez et al. 2015] Alvarez, R.; Garcia, D.; Moreno, Y.; and Schweitzer, F. 2015. Sentiment cascades in the 15m movement. *EPJ Data Science* 4(1):1–13.
- [Aral and Van Alstyne 2011] Aral, S., and Van Alstyne, M. W. 2011. The Diversity-Bandwidth tradeoff. *American Journal of Sociology* 117(1):90–171.
- [Backstrom, Sun, and Marlow 2010] Backstrom, L.; Sun, E.; and Marlow, C. 2010. Find me if you can: improving geographical prediction with social and spatial proximity. In *WWW2010*, 61–70.
- [Bakshy et al. 2012] Bakshy, E.; Rosenn, I.; Marlow, C.; and Adamic, L. 2012. The role of social networks in information diffusion. In *WWW2012*.
- [Barthlemy 2011] Barthlemy, M. 2011. Spatial networks. *Physics Reports* 499(1–3):1–101.
- [Brashears 2013] Brashears, M. E. 2013. Humans use compression heuristics to improve the recall of social networks. *Nature Scientific Reports* 3.
- [Burt 1995] Burt, R. 1995. *Structural Holes: The Social Structure of Competition*. Cambridge, MA: Harvard University Press.
- [Burt 2004] Burt, R. S. 2004. Structural holes and good ideas. *The American Journal of Sociology* 110(2):349–399.
- [Cheng et al. 2011] Cheng, Z.; Caverlee, J.; Lee, K.; and Sui, D. 2011. Exploring millions of footprints in location sharing services. In *ICWSM2011*.
- [Cramer, Rost, and Holmquist 2011] Cramer, H.; Rost, M.; and Holmquist, L. E. 2011. Performing a check-in: Emerging practices, norms and 'conflicts' in location-sharing using foursquare. In *Proc. 13th International Conference on Human Computer Interaction with Mobile Devices and Services*.
- [Duffy 2011] Duffy, S. 2011. A community creating their own rules on foursquare. In *Proc. First NSS Workshop on Mobile and Online Social Networks*.
- [Dunbar 2003] Dunbar, R. 2003. Evolution of the social brain. *Science* 302(5648):1160–1161.
- [Eagle, Macy, and Claxton 2010] Eagle, N.; Macy, M.; and Claxton, R. 2010. Network diversity and economic development. *Science* 328(5981):1029–1031.
- [Fontaine et al. 2007] Fontaine, J. R.; Scherer, K. R.; Roesch, E. B.; and Ellsworth, P. C. 2007. The world of emotions is not two-dimensional. *Psychological science* 18(12):1050–1057.
- [Fredrickson 2001] Fredrickson, B. L. 2001. The role of positive emotions in positive psychology: The broaden-and-build theory of positive emotions. *American psychologist* 56(3):218.
- [Garas et al. 2012] Garas, A.; Garcia, D.; Skowron, M.; and Schweitzer, F. 2012. Emotional persistence in online chatting communities. *Scientific Reports* 2.
- [Garcia and Schweitzer 2011] Garcia, D., and Schweitzer, F. 2011. Emotions in product reviews—empirics and models. In *SocialCom2011*.
- [Garcia et al. 2012] Garcia, D.; Mendez, F.; Serdült, U.; and Schweitzer, F. 2012. Political polarization and popularity in online participatory media: an integrated approach. In *Proc. first edition workshop on Politics, elections and data*.
- [Garcia, Garas, and Schweitzer 2012] Garcia, D.; Garas, A.; and Schweitzer, F. 2012. Positive words carry less information than negative words. *EPJ Data Science* 1:3.
- [Garcia-Gavilanes, Quercia, and Jaimes 2013] Garcia-Gavilanes, R.; Quercia, D.; and Jaimes, A. 2013. Cultural dimensions in twitter: Time, individualism and power. In *ICWSM2013*.
- [Golder and Macy 2014] Golder, S. A., and Macy, M. W. 2014. Digital footprints: opportunities and challenges for online social research. *Sociology* 40(1):129.
- [Gonçalves et al. 2013] Gonçalves, P.; Araújo, M.; Benvenuto, F.; and Cha, M. 2013. Comparing and combining sentiment analysis methods. In *COSN*.
- [Gonçalves, Perra, and Vespignani 2011] Gonçalves, B.; Perra, N.; and Vespignani, A. 2011. Modeling users activity on twitter networks: Validation of dunbars number. *PLoS one* 6(8):e22656.
- [González-Bailón, Banchs, and Kaltenbrunner 2012] González-Bailón, S.; Banchs, R. E.; and Kaltenbrunner, A. 2012. Emotions, public opinion, and us presidential approval rates: A 5-year analysis of online political discussions. *Human Communication Research* 38(2):121–143.
- [Granovetter 1973] Granovetter, M. 1973. The Strength of Weak Ties. *The American Journal of Sociology* 78(6):1360–1380.
- [Granovetter 1983] Granovetter, M. 1983. The strength of weak ties: A network theory revisited. *Sociological Theory* 1:201–233.
- [Hofstede 1980] Hofstede, G. 1980. *Culture's Consequences: International Differences in Work-Related Values*, volume 1980. SAGE.
- [Kang and Lerman 2015] Kang, J., and Lerman, K. 2015. User effort and network structure mediate access to information in networks. In *ICWSM2015*.
- [Kayes et al. 2015] Kayes, I.; Kourtellis, N.; Quercia, D.; Iamnitchi, A.; and Bonchi, F. 2015. Cultures in community question answering. In *Hypertext*, 175–184. ACM.
- [Krackhardt 1987] Krackhardt, D. 1987. Cognitive social structures. *Social networks* 9(2):109–134.
- [Kucuktunc et al. 2012] Kucuktunc, O.; Cambazoglu, B. B.; Weber, I.; and Ferhatosmanoglu, H. 2012. A large-scale sentiment analysis for yahoo! answers. In *WSDM2012*.
- [Liben-Nowell et al. 2005] Liben-Nowell, D.; Novak, J.; Kumar, R.; Raghavan, P.; and Tomkins, A. 2005. Geographic routing in social networks. *PNAS* 102(33):11623–11628.
- [Mitchell et al. 2013] Mitchell, L.; Frank, M. R.; Harris, K. D.; Dodds, P. S.; and Danforth, C. M. 2013. The geography of happiness: Connecting twitter sentiment and expression, demographics, and objective characteristics of place. *PLoS ONE* 8(5):e64417.

- [Niedenthal and Brauer 2012] Niedenthal, P. M., and Brauer, M. 2012. Social functionality of human emotion. *Annual review of psychology* 63:259–285.
- [Noulas et al. 2011] Noulas, A.; Scellato, S.; Mascolo, C.; and Pontil, M. 2011. An empirical study of geographic user activity patterns in foursquare. In *ICWSM2011*.
- [Onnela et al. 2007] Onnela, J.; Saramäki, J.; Hyvönen, J.; Szabó, G.; Lazer, D.; Kaski, K.; Kertész, J.; and Barabási, A. 2007. Structure and tie strengths in mobile communication networks. *PNAS* 104(18):7332–7336.
- [Osgood, Suci, and Tannenbaum 1964] Osgood, C. E.; Suci, G. J.; and Tannenbaum, P. H. 1964. *The measurement of meaning*. University of Illinois Press.
- [Pang and Lee 2008] Pang, B., and Lee, L. 2008. Opinion mining and sentiment analysis. *Foundations and trends in information retrieval* 2(1-2):1–135.
- [Putnam 2000] Putnam, R. D. 2000. *Bowling Alone: The Collapse and Revival of American Community*. Simon & Schuster.
- [Quercia et al. 2012] Quercia, D.; Ellis, J.; Capra, L.; and Crowcroft, J. 2012. Tracking “gross community happiness” from tweets. In *Proceedings of the CSCW’12 Conference*.
- [Quercia, Capra, and Crowcroft 2012] Quercia, D.; Capra, L.; and Crowcroft, J. 2012. The social world of twitter: Topics, geography, and emotions. In *ICWSM2-12*.
- [Rajadesingan, Zafarani, and Liu 2015] Rajadesingan, A.; Zafarani, R.; and Liu, H. 2015. Sarcasm detection on twitter: A behavioral modeling approach. In *WSDM2015*.
- [Rimé 2009] Rimé, B. 2009. Emotion elicits the social sharing of emotion: Theory and empirical review. *Emotion Review* 1(1):60–85.
- [Russell and Mehrabian 1977] Russell, J. A., and Mehrabian, A. 1977. Evidence for a three-factor theory of emotions. *Journal of research in Personality* 11(3):273–294.
- [Sampson, Raudenbush, and Earls 1997] Sampson, R. J.; Raudenbush, S. W.; and Earls, F. 1997. Neighborhoods and violent crime: A multilevel study of collective efficacy. *Science* 277(5328):918–924.
- [Shea et al. 2015] Shea, C. T.; Menon, T.; Smith, E. B.; and Emich, K. 2015. The affective antecedents of cognitive social network activation. *Social Networks* 43:91–99.
- [Thelwall et al. 2013] Thelwall, M.; Buckley, K.; Paltoglou, G.; Skowron, M.; Garcia, D.; Gobron, S.; Ahn, J.; Kappas, A.; Küster, D.; and Holyst, J. A. 2013. Damping sentiment analysis in online communication: discussions, monologs and dialogs. In *Computational linguistics and intelligent text processing*. Springer. 1–12.
- [Thelwall, Buckley, and Paltoglou 2012] Thelwall, M.; Buckley, K.; and Paltoglou, G. 2012. Sentiment strength detection for the social web. *Journal of the American Society for Information Science and Technology* 63(1):163–173.
- [Travers and Milgram 1969] Travers, J., and Milgram, S. 1969. An experimental study of the small world problem. *Sociometry* 32(4):425–443.
- [Tufekci 2014] Tufekci, Z. 2014. Big questions for social media big data: Representativeness, validity and other methodological pitfalls. In *ICWSM2014*.
- [Warriner, Kuperman, and Brysbaert 2013] Warriner, A. B.; Kuperman, V.; and Brysbaert, M. 2013. Norms of valence, arousal, and dominance for 13,915 english lemmas. *Behavior research methods* 45(4):1191–1207.
- [Watson, Clark, and Tellegen 2013] Watson, D.; Clark, L. A.; and Tellegen, A. 2013. Development and validation of brief measures of positive and negative affect: the panas scales. *Journal of personality and social psychology* 3(6):1063.
- [Webb et al. 1966] Webb, E. J.; Campbell, D. T.; Schwartz, R. D.; and Sechrest, L. 1966. *Unobtrusive measures: Non-reactive research in the social sciences*, volume 111. Rand McNally Chicago.