Original Research

Leveraging geotagged Twitter data to examine neighborhood happiness, diet, and physical activity

AUTHORS: Quynh C Nguyen ^a, Suraj Kath ^b, Hsien-Wen Meng ^a, Dapeng Li ^c, Ken Robert Smith ^d, James A. VanDerslice ^e, Ming Wen ^f, Feifei Li ^b

AFFILIATIONS: ^aDepartment of Health Promotion and Education, College of Health, University of Utah, Salt Lake City, UT, USA; ^bSchool of Computing, University of Utah; ^cDepartment of Geography, University of Utah; ^dDepartment of Family and Consumer Studies, University of Utah; ^eDivision of Family and Preventive Medicine, School of Medicine, University of Utah; ^fDepartment of Sociology, University of Utah

CORRESPONDENCE TO: Quynh Nguyen, Department of Health Promotion and Education, University of Utah, 1901 E South Campus Drive, Annex B 2124, Salt Lake City, UT, 84112

Tel: (801) 585-5134

Email: quynh.nguyen@health.utah.edu

Manuscript word count: 5548

Abstract word count: 178

Tables: 6

Figures: 3

ABSTRACT

Objectives: Using publicly available, geotagged Twitter data, we created neighborhood indicators for happiness, food and physical activity for three large counties: Salt Lake, San Francisco and New York. **Methods:** We utilize 2.8 million tweets collected between February-August 2015 in our analysis. Geo-coordinates of where tweets were sent allow us to spatially join them to 2010 census tract locations. We implemented quality control checks and tested associations between Twitter-derived variables and sociodemographic characteristics. **Results:** For a random subset of tweets, manually labeled tweets and algorithm labeled tweets had excellent levels of agreement: 73% for happiness; 83% for food, and 85% for physical activity. Happy tweets, healthy food references, and physical activity references were less frequent in census tracts with greater economic disadvantage and higher proportions of racial/ethnic minorities and youths. **Conclusions:** Social media can be leveraged to provide greater understanding of the well-being and health behaviors of communities—information that has been previously difficult and expensive to obtain consistently across geographies. More open access neighborhood data can enable better design of programs and policies addressing social determinants of health.

Keywords: Twitter messaging; neighborhood; happiness; physical activity; diet; food

1. INTRODUCTON

The literature examining neighborhood effects on health has flourished in the last decade.¹ Extant research has provided evidence on associations between the neighborhood environment and mortality risk²⁻⁸, life expectancy⁹, mental health¹⁰, self-rated health¹¹, obesity,¹²⁻¹⁵, and diabetes^{16,17}—even after adjusting for individual characteristics. Poor access to healthy food¹⁸⁻²³, fast food chains²⁴, the lack of recreational facilities^{25,26}, and higher crime rates^{19,27} all correlate with higher obesity rates. Community happiness levels also have been inversely related to obesity as well as other outcomes including hypertension, suicide, and life expectancy.²⁸⁻³³ Adverse neighborhood conditions concentrate in poor, minority neighborhoods^{13,34-36}, thereby increasing health disparities. Furthermore, the epidemic rise in obesity and related chronic diseases in recent decades signal the importance of structural forces and social processes.

Nonetheless, the dearth of data on contextual factors limits the investigation of multilevel effects on health. Certain places^{37,38} have extensive neighborhood data collected on them, but they are the exception rather than the rule, and it is difficult to make comparisons across geographies because available measures vary greatly across them. Also patterns seen in specific places may not apply to other places. For instance, estimates and patterns seen in urban areas may not apply to rural areas. Neighborhood data collection is expensive and time consuming, and then only available for certain places or time periods and become outdated quickly.³⁹ Moreover, while comparable neighborhood data across large areas are highly lacking, the neighborhood data we do have are typically data on compositional characteristics (e.g., percent females) and features of the built environment (e.g., number of grocery stores and health care clinics). These data do not capture the social environment, or an individual's interactions with that environment.

Social processes and networks can affect health via a myriad of mechanisms, such as 1) the maintenance of norms around healthy behaviors via informal social control; 2) the stimulation of new interests such as a new sport or exercise; 3) political advocacy for access to neighborhood amenities and protection against stressors and toxic agents; 4) emotional support; and 5) the dispersal of knowledge about health promotion practices.⁴⁰⁻⁴⁴ According to Social Learning Theory, learning takes place in a social context.⁴⁵ Behaviors are adopted by observing how the behavior is performed by others, attitudes around that behavior, and outcomes associated with that behavior. Empirically, the adoption of specific health behaviors related to food consumption, health screening, smoking, alcohol consumption, drug use, and sleep has been observed to disperse through social networks.⁴⁶⁻⁵¹ Similarly, evidence suggests that emotional states such as mood⁵², happiness⁵³, depression⁵⁴, and suicidality⁵⁵ can spread through social networks. The measurement of area-level happiness and subjective-well-being is a new and expanding research endeavor.⁵⁶⁻⁶⁰ For instance, in 2012, the United Nations began its annual release of a World Happiness Report.⁵⁶ Social media may influence individuals' health behaviors but may also be a way to characterize prevalent community characteristics and patterns of behaviors.

1.1 Study aims

Given the vast literature documenting the influence of social networks on individual health behaviors and health outcomes, we believe that social media data represent an important new data resource for neighborhood researchers. Thus, using publicly available, geotagged Twitter data, we construct novel indicators of neighborhood happiness levels, healthiness of food, and physical activity. We conduct quality control activities and perform validation analysis comparing Twitter-derived neighborhood indicators to demographic and economic

characteristics of the corresponding census tract. In order to test our computer algorithm for constructing neighborhood indicators, we selected three counties that display diversity in regards to geographical location, landscape, housing market, cultural characteristics, and demographic characteristics (e.g., racial/ethnic composition, age distribution, and household size). The three counties are the following: Salt Lake County, San Francisco County, and New York County.

2. METHODS

2.1 Social media data collection

From February – August 2015, we utilized Twitter’s Streaming Application Programming Interface (API) to continuously collect a random 1% subset of publicly available tweets with latitudes and longitude coordinates. We present in-depth analyses and findings for three counties in the United States: Salt Lake County (367,204 tweets); San Francisco County (same as San Francisco city; 653,670 tweets); and New York County (1,828,026 tweets).

2.2 Spatial join

We linked 99.8% of tweets with available GPS coordinates to their respective 2010 census tract locations. We used Python and relevant GIS libraries (Shapely and Fiona) to accomplish this task. An R-Tree was used to build a spatial index⁶¹ on census tract polygon data. R-tree indexing allows for faster spatial searches on the data because the R-tree groups data into bounding rectangles and narrows the search space. A query that does not intersect the bounding rectangle cannot intersect any of its component parts. Utilizing the latitude and longitude coordinates of where tweets were sent, spatial joins were performed on tweets to identify the corresponding census tracts. Tweets that were not assigned a census location included those with destinations bordering the United States (i.e., Mexico and Canada).

2.3 Processing tweets

We processed tweets to create variables that measure sentiment, food, and physical activity. To accomplish this task, we utilized a bag-of-words algorithm which creates a simplifying representation of tweets that disregards grammar and word order, but has the capacity to track the frequency of terms or components of tweets, and then performs computations on those components and terms. Several steps were conducted that first included dividing each tweet into tokens.⁶² A tokenizer divides text into a sequence of tokens, which roughly correspond to “words.” We are using a tokenizer particularly suitable for processing English text called the PTBTokenizer (aka the Stanford Tokenizer). The PTBTokenizer is an efficient, fast, and deterministic tokenizer. It can tokenize text at a rate of about 1,000,000 tokens per second on a standard personal computer. Utilizing heuristics, it can usually differentiate when single quotes are part of words and when periods do and do not imply sentence boundaries. After we obtain the tokens (i.e., individual words) from a tweet, we then search each word in our word dictionary to get its corresponding happiness score for sentiment analysis. Currently, we are ignoring words which are not present in our word dictionary. Using this algorithm, sentiment scores can be assigned to approximately 80-85% of tweets across geographies. Below we describe in more detail our algorithms for constructing each variable.

2.4 Sentiment analysis

To implement sentiment analysis, we utilized the Language Assessment by Mechanical Turk (LabMT) word list, compiled by obtaining the most frequently occurring words in each of the following four text sources: Google Books (English), music lyrics, the New York Times and Twitter. About 10,000 of these individual words have been scored by users of Amazon’s

Mechanical Turk service on a scale of 1 (sad) to 9 (happy), resulting in a measure of average happiness for each given word. We extended the LabMT list by adding 158 commonly used emoticons – pictorial representation of a feeling or facial expression expressed through a combination of numbers, letters, and punctuation marks. For instance, :-) represents a smiley face. Happiness values for the emoticons were assigned according to the word meaning of the emoticons.

Average happiness for each tweet was computed by taking the average of happiness values of words composing the tweet, excluding neutral words (e.g., “the,” “is”).⁶³ The happiness value for a census tract was then computed as the average of sentiment values of all tweets in that census tract. Our algorithm tracked and accounted for intensifiers (such as the use of exclamation marks and all-capitalized words) in computing sentiment. For tweets with happiness scores above 6, we added 1 for the use of one exclamation mark (7.2% of tweets) and 2 for use of 2 exclamation marks (4.7% of tweets). For tweets with happiness scores less than 4, 1 and 2 was subtracted, respectively. For tweets with at least two words in which all the letters were capitalized (5.7% of tweets), 1 was added to the happiness score of tweets with scores originally above 6 and subtracted from the happiness scores of tweets with original scores below 4. If three or more words were capitalized, 2 was added if the happiness score was originally above 6 and 2 was subtracted if the happiness score was below 4. We did not modify tweets with sentiment scores between 4 and 6 (i.e., neutral tweets) because the directionality of the emphasis was ambiguous. We required that at least two words be capitalized in order to avoid extra weight being given to tweets that utilized capitalized state initials (e.g., CA) or denotations of time (e.g., “AM” or “PM”). Tweets can have use of exclamation marks and at least two all-caps words (1.4% of tweets).

For quality control, we evaluated our sentiment analysis against the ratings of human coders.⁶⁴ We created 20 online surveys consisting of 25 publicly available tweets each (500 tweets in total). The 500 tweets were randomly selected within sentiment groups such that 50% of tweets were happy, 25% were sad and 25% were neutral. We randomly sorted this master list of tweets and created 20 online surveys from the sorted list. Each online survey received 15 responses. Upon completion of the survey (i.e., participant rated all 25 tweets), 25 cents (\$0.25) was deposited into the participant's Amazon Mturk account per online survey completed. The study was approved by the University of Utah Institutional Review Board. A total of 32 people participated in our study (participants generally did several batches of tweets rather than just one batch). Study time was four days (April 1 – April 5, 2015). We utilized the human ratings to calibrate the scaling of our sentiment analysis. The bag of words technique utilized by Mitchell and colleagues⁶³ categorized happy, sad and neutral tweets as having the following sentiment scores: ≥ 6 , ≤ 4 , and 4-6. However, after calibration to human ratings such that the proportion of happy, sad and neutral tweets matched human categorization, sentiment scores were categorized as such ≥ 6.6 (30%), ≤ 4.9 (33%), and 4.9-6.6 (37%). The kappa statistics for agreement between human vs. algorithm-derived sentiment ratings was 73%, if we dichotomize sentiment as happy vs. not happy at the modified cut points. Additionally, sensitivity analyses with a different set of weights for capitalizations and exclamation marks (e.g., 0.5 and 1.0, 1.0 and 2.0, and 2.0 and 3.0) resulted in qualitatively similar agreement between human coders and the algorithm (70-71%) among tweets with intensifiers.

2.5 Food analysis

We compiled a list of over 1,300 popular food words from the U.S. Department of Agriculture's National Nutrient Database.^{65,66} Each food item was associated with a measure of

caloric density, operationalized as calories per 100 grams. Fruits, vegetables, nuts, and lean proteins (e.g., fish, chicken, and turkey) were additionally labeled as “healthy foods.” We excluded fried foods from our count of healthy foods. Our food list also contained 48 popular national fast food restaurants such as McDonald’s and Kentucky Fried Chicken to enable quantification of fast food references.

To analyze food culture, each tweet was examined for words or phrases matching those on our list. Each food item on our list was described by one or two words. Our computer algorithm first searched over a tweet for matches to two-word foods (e.g., orange chicken). It then searched over the remaining words for matches to one-word food terms (e.g., taco). We computed caloric density by summing up all the foods mentioned in the tweet. We also created a count of healthy food references and fast food restaurant references for each tweet.

We attempted to calculate happiness scores for all tweets, including those involving food, using the algorithm described in our sentiment analysis section. For food tweets, separate variables were calculated for the average sentiment of tweets involving healthy foods and fast food. That is, if a food tweet mentioned a fast food restaurant, its happiness score was included in our calculation of the average sentiment around fast food. Alternatively, if a food tweet mentioned a healthy food, it was included in our calculation of the average sentiment around healthy foods. These variables (i.e., any food references, healthy food references, fast food references, caloric density, and sentiment towards healthy foods and fast food) were then aggregated and summarized at the census-tract level to create neighborhood indicators of food culture.

Quality control activities were implemented to determine whether tweets identified by our algorithm as food-related would be rated as pertaining to food by humans. Two of the co-authors rated a random assortment of 2500 tweets. Initial inter-rater reliability was 92% and discordant values were reviewed until 100% agreement between raters was reached. For tweets that our algorithm had labeled as non-food related, 81% of those labels were correct according to human ratings. For tweets that our algorithm labeled as food related, 83% of those labels were correct. The kappa statistics for agreement between labels produced by our algorithm and human coders was 83%. Our algorithm missed food references if the food term was not included in our food list. We excluded food items with many non-related food meanings such as “perch.” For tweets that had been mislabeled as food-related, common reasons included the following: food term utilized in the tweet had a non-food related meaning (e.g., pho fighters); the term was utilized in a food pun or metaphor; or the tweet was a food advertisement.

2.6 Physical activity/recreation analysis

We created a list of physical activities using published lists of physical activity terms gathered from physical activity questionnaires, compendia of physical activities, and popularly available fitness programs.^{67,68} Our physical activity list has 258 different activities that incorporate gym-related exercise (e.g., treadmill, weight lifting), sports (e.g., baseball), recreation (e.g., hiking, scuba diving) and household chores (e.g., gardening). Using Metabolic Equivalents (METs) associated with physical activities, we quantified the caloric expenditure of each physical activity mention, scaled for a duration of 30 minutes and for a 155 pound individual.⁶⁹ Upon piloting our algorithm, we identified commonly utilized phrases or pop culture references that do not involve physical activity (i.e., walking dead, running late, yoga pants) which were manually coded and excluded. Moreover, in order to help reduce the

possibility that the tweet was about watching rather than actually participating in the physical activity, we excluded the tweet if it contained any of the following terms: watch, watching, watches, watched, attend, attending, attends, and attended. In reviewing preliminary labeled physical activity data, we found that most tweets (over 90%) pertaining to team sports (i.e., baseball, basketball, football, soccer) were about watching games rather than participating in them. Thus, for team sports, we also required that the tweet include the words play, playing, or played. Our algorithm created the following physical activity variables for each tweet: indicator variable for any physical activity mention, indicator variable for gym-related vs. other types of physical activity, sum of caloric expenditure of mentioned activities, and sentiment around physical activity.

Quality control activities were implemented to determine accuracy of tweets that had been labeled by our algorithm as being related to physical activity. Two annotators labeled a random assortment of 2500 tweets with an initial kappa agreement of 95%. Again discordant values were reviewed until 100% inter-rater agreement was reached. For tweets that the computer algorithm had labeled as not related to physical activity, 97% of those labels were correct according to human ratings. For tweets that our algorithm labeled as related to physical activity, 82% of those labels were correct. Kappa agreement between labels generated by our algorithm and human coders was 85%. Common classification errors included the following: 1) tweet was about watching a professional sport rather than participating in it and 2) tweet utilized a term that took on a non-physical activity related meaning (e.g., running a fever).

2.7 Census tract characteristics

To examine how Twitter-derived neighborhood variables relate to more traditional neighborhood variables, we merged our social media dataset with the 2010 Census and 2013 American Community Survey data which comprised the following demographic, household and economic characteristics: percent 65 years+; percent 10-24 years; percent male; percent African American; percent white; percent Hispanic; percent in households with relatives (other than spouse and children); percent in households with unmarried partner; household size; percent single female-headed households; percent householder living alone; percent owner-occupied housing; median family income; percent college graduates; percent unemployed; percent with less than a high school degree; percent families living in poverty (Table 1).

2.8 Analytic approach

We grouped tweets by census tract and created indicators of happiness, food, and physical activity for each census tract. We then merged our tract-level social media database with census data (see above). To facilitate interpretation of findings for different variables, we standardized all variables to have a mean of zero and standard deviation of one. We implemented linear regression models to examine associations between Twitter-derived neighborhood variables and census tract characteristics for three demonstrative counties: Salt Lake, San Francisco, and New York. In addition, we mapped the distribution of happiness scores across census tracts in the three counties. We ranked food and physical activity terms by popularity, and reported the most popular terms.

3. RESULTS

Across the three counties, tweets were more likely to be neutral or positive rather than negative in sentiment (Table 1). Prevalence of happy tweets was highest in New York, followed

by San Francisco and Salt Lake counties. About 3.1-6.6% of the tweets were food-related. Of these food tweets, about 16-17% mentioned healthy foods and 6-10% mentioned fast food restaurants. The mean (standard deviation) caloric density of food references was 250-261 calories (per 100 grams). Food tweets that did not include fast food were slightly happier than those with fast food references ($p < 0.01$). About 1.7% of tweets were about physical activity, and most physical activity references were not gym-related (e.g., skiing, hiking)—although higher proportions of gym-related activities were reported in New York than San Francisco and Salt Lake counties. The mean caloric expenditure of physical activity references was 217-219 calories (assuming duration of 30 minutes for a 155 pound person). Positive sentiment around physical activity was higher in New York compared to Salt Lake County ($p < 0.001$).

The demographic and economic composition of residents differed in the three counties. More youth were present in Salt Lake County than the other two (Table 1). Salt Lake County had fewer racial/ethnic minorities and larger household sizes. Median family incomes and percent college graduates were much higher in New York and San Francisco than in Salt Lake County. New York had a higher poverty rate than Salt Lake or San Francisco.

While possible happiness scores can range from 1 to 9, at the census tract level, most locations had scores that varied between 6.0 and 7.0 (Figures 1-3). The majority of areas in Salt Lake County had sentiment scores between 6.0-6.5, although lower sentiment scores were more prevalent on the west and south sides (Figure 1). For San Francisco, happiness scores were highest near areas bordering the bay or ocean (e.g., Presidio, Embarcadero) as well as areas near Golden Gate Park and the Mission District (Figure 2). Of the three counties, the proportion of tweets characterized as happy was highest in New York County and hotspots of happiness were

found near Central Park, Meatpacking District, Greenwich Village, Tribeca, Soho and also Wards Island and upper areas like Inwood (Figure 3).

We conducted crude analyses in which each demographic and economic characteristic was used as a regressor for standardized tract level happiness. Higher proportions of whites and college graduates were positively associated with higher tract happiness levels (eTable 1). Conversely, higher proportion of young people aged 10-24 years, Hispanics, single female-headed households, and higher unemployment rate, poverty rate, and household sizes were associated with lower happiness.

We additionally implemented adjusted analyses controlling for all predictors simultaneously. However, due to high collinearity of census tract characteristics, we reduced the number of variables included in the adjusted models and factor-analyzed four highly correlated characteristics (i.e., percent female-headed households, percent families living in poverty, unemployment rate, and median family income) to create an economic disadvantage factor score. In San Francisco, percent African American and greater household size were related to lower happiness scores (Table 2). In Salt Lake County, higher percentages of elderly individuals and males were associated with higher happiness scores while greater household size, percent Hispanic, and economic disadvantage were related to lower happiness scores—although the latter two characteristics bordered statistical significance (Table 2). In New York City, proportion of young people aged 10-24 years was linked to lower happiness scores. Sensitivity analyses using dichotomized measures of sentiment (i.e., happy vs. not happy; sad vs. not sad) resulted in finding the same statistically significant associations with census tract characteristics as compared to analyses utilizing continuous happiness scores (not shown).

Although our food list consisted of over 1300 food terms, 50% of food tweets could be characterized by around 25 of the most popular food terms (Table 3). Across the three counties, top food terms included coffee, beer, tea, wine, pizza, burger, ice cream, chocolate, cake, chicken, sushi, and salad among others. In crude analyses, healthy food references increased with higher proportions of unmarried cohabitating adults, householder living alone, and college graduates (eTable 2). There is also some suggestion that healthy foods references increase with higher proportion of whites. Healthy food references were less frequent in tracts with higher proportions of individuals aged 10-24 years, larger household sizes, and more economic disadvantage.

In adjusted analyses, increasing household size and greater proportions of individuals aged 10-24 were associated with less frequent healthy food references (Table 4). In San Francisco, we also saw that greater percent of elderly, males, and African Americans were related to fewer healthy food references (Table 4). No sociodemographic characteristics statistically significantly predicted sentiment around healthy foods (not shown). In New York County, greater household size ($B = -0.43$; $p = 0.02$) and economic disadvantage ($B = -0.45$; $p = 0.04$) predicted less positive sentiment around fast food. Additionally, in New York ($B = -0.55$; $p < 0.002$) economic disadvantage was linked to lower frequency of fast food tweets. However, in Salt Lake County, economic disadvantage was marginally related to more fast food tweets ($B = 0.24$; $p = 0.09$).

Although our physical activities list comprised more than 250 terms, across the three counties, 75% of physical activity references could be characterized by about 10 of the most popular physical activity terms. Top physical activity terms included the following: walking, running, dance, golf, hiking, skiing, yoga, and workout (Table 5). In crude analyses, predictors of

tweets involving physical activity were similar between San Francisco and Salt Lake counties – with percent aged 10-24 years, percent Hispanics, and percent in households with relatives associated with lower frequency of physical activity tweets. In San Francisco and Salt Lake counties, percent college graduates was associated with higher frequency of physical activity tweets. In New York County higher median income was associated with higher frequency of physical activity tweets (eTable 3). In adjusted models for physical activity, higher proportions of individuals aged 65+ were associated with higher frequency of physical activity tweets in San Francisco. In Salt Lake and New York counties, economic disadvantage was related to lower frequency of physical activity tweets (Table 6). Larger household size was associated with lower positive sentiment around physical activity in San Francisco ($B=-0.21$, $p=0.02$) and Salt Lake ($B=-0.24$; $p=0.05$) counties. We also examined predictors of caloric expenditure of activities mentioned in tweets; higher percent of males (New York: $B=-0.27$; $p=0.01$; San Francisco: $B=-0.15$; $p=0.10$) and economic disadvantage (New York: $B=-0.49$; $p=0.005$; San Francisco: $B=-0.23$; $p=0.10$) were related to lower caloric expenditure. Additionally, in New York greater household size ($B=0.60$; $p=0.001$) was related to higher caloric expenditure.

DISCUSSION

Social media is a massive data resource that is beginning to be leveraged for health research such as the prediction of flu⁷⁰, predicting the onset of depression⁷¹, modeling outbreaks,⁷² characterization of emergency response⁷³, investigating spatial patterns in obesity-related tweets and their proximity to McDonalds⁷⁴, and identifying associations between healthful tweets and presence of green retailers.⁷⁵ In this study, across three different counties in the United States, we generally find that happy tweets, healthy food references, and physical activity references were less frequent in census tracts with greater economic disadvantage and

higher proportions of racial/ethnic minorities and youths. However, we did not find consistent relationships between economic disadvantage and fast food tweets.

Nascent work suggests that social media interactions may serve to share knowledge, provide support, and influence behaviors and mental states including happiness.⁷⁶⁻⁷⁸ Our algorithms demonstrated levels of accuracy comparable to other natural language processing techniques.^{79,80} Nonetheless, while sentiment analysis is an emerging endeavor in computer science and natural language processing, very few studies have been conducted on the classification of food and physical activity references using social media. Additionally, our algorithms are unique in creating indicators at the neighborhood level. The vast majority of Twitter analyses are conducted at the city, county or state levels. We plan to make our database, *HashtagHealth*, available to other public health researchers upon one full year of data collection.

Our sentiment analysis was adapted from methodology developed by Dodds et al.³² and further tested by Mitchell et al.⁶³ Our sentiment analysis extended the methodology by accounting for emoticons and use of capitalizations and exclamation marks in estimating the sentiment of a tweet. Mitchell and colleagues found that higher socioeconomic status was associated with higher happiness scores at the city level. Moreover, they found a mild ($r=-0.34$) correlations between happiness and obesity rates for 190 Metropolitan Statistical Areas (MSAs)⁶³ and that happiness scores moderately correlated with other state-level indicators of well-being including shootings, the Peace index, America's Health Ranking, and the Gallup-Healthways Well-Being Index (correlations ranged between 0.51 and 0.64).⁶³ Widener and Li conducted Twitter analysis of food references and found that disadvantaged areas had fewer positive references for fruits and vegetables.⁸¹ Ghosh and Guha found a positive correlation between obesity prevention themed tweets and number of policies related to obesity, nutrition

and physical activity at the state level.⁷⁴ They also found a strong positive correlation between tweets about high calorie foods/obesity and locations of McDonalds.⁷⁴ Zhang and colleagues manually coded a random assortment of 30,000 tweets about physical activity, finding that most were neutral in sentiment and only 9.0% offered social support.⁶⁷

4.1 Study strengths and limitations

Our Twitter database leverages publicly exhibited text to construct indicators of the social environment previously not available to neighborhood effects researchers. This project is made possible through collaborations between public health, social sciences, and computer science. Using social media for neighborhood research purposes has substantial strengths including the public nature of tweets, the regularity and pervasiveness of tweets, which enables easy updating of constructed indicators, and participation from increasing segments of the population. Although we utilized Twitter as our main source of social media data, Twitter users with multiple social media accounts such as Facebook, Instagram and Vine will often link their accounts to display their posts on multiple platforms. We found tweets often incorporated links to Facebook, Instagram and Vine posts.

Nonetheless, social media as a data resource is not without limitations which include an over-representation of young individuals. While more users of social media tend to be younger, adoption rates have been steadily increasing over the years for all age groups and as of 2014, usage of social networking sites ranged from 89% among 18-29 year olds to 49% among those 65 years and older.⁸² Moreover, arguably a person can be affected by social media even if he or she is not an active user, for instance, through discussions by friends and colleagues or headlines in the news that now regularly feature social media content. Tweets also include information

rarely found in other neighborhood sources. Twitter users are composed of individuals as well as groups of individuals, organizations, companies and news outlets. Thus, compiling such information may allow for a more comprehensive examination of the social environment.

Nevertheless, social media researchers can only harvest and utilize information that people are willing to share. This information includes their activities, intentions, and opinions. Also, our construction of neighborhood indicators from Twitter data necessitated that we restrict our data collection to geotagged tweets (i.e., tweets in which users enabled location on their mobile phones). Previous studies suggest that about 1-2% of tweets may contain GPS location information, and that use of Twitter's Streaming API may obtain 90% of all geotagged tweets (in comparison to Twitter's Firehouse which returns all tweets).^{83,84} Tweets with location information may differ from those who do not. For instance, our body of tweets may have a higher proportion of public and social activities such as friends tweeting from a restaurant or an event.

Moreover, tweets were spatially joined to the location of where they were sent, not the user profile location (which is a "text entry" of home city, state, or country). Thus, tweets aggregated to a census tract can be composed of those written by people visiting the area and those living there. We believe that including data from visitors is a strength of the study design because people can be mobile throughout the course of the day and regardless of their residency⁸⁵, people can directly influence the social environment in which they are in physical proximity. Additionally some places such as New York City's Times Square and Fisherman's Wharf are characterized by the large influx of visitors. A potential concern is that the demographic or economic characteristics of visitors may differ from the compositional characteristics of residents. However, people may tend to visit areas that are aligned with their

own background characteristics (i.e., assortative sorting) as suggested by the literature on neighborhood segregation.^{86,87} Moreover, the data is naturally weighted towards residents or people who spend greater amounts of time in a census tract (given that we are collecting a random subset of tweets and people who spend more time in a census tract are more likely to have their tweet included in our dataset). Nonetheless, the influence of a tweet can, of course, extend much further than the actual physical location of where the tweet was sent.

An additional complexity is that the location of where the tweet was sent may differ from where the sentiment, event or activity originated. For instance, frustrations about work could be expressed later in the day via a tweet at a bar. Thus having only access to the geographic information of where the tweet sent may mean that we are obtaining only partial information of the all possible geographies described in the tweet. This potential “migration bias” may reduce the strength of observed associations.

Moreover, taking census tracts as the neighborhood unit can be problematic given the variation in population density and physical size. However, census tracts are useful neighborhood approximations given that they are designed to be relatively homogenous with respect to sociodemographic composition and living conditions. Additionally, publicly available administrative data is available at the census tract level. Aggregating tweets to the census tract level also ameliorates concerns about privacy because individual-level tweets and any identifying information are not be made available. Our neighborhood dataset contains only summaries of group-level characteristics. In addition, geotagged data were voluntarily contributed by Twitter users via an opt-in mechanism on the Twitter mobile app.

In creating our neighborhood indicators from Twitter data, we prioritized transparency and ease of implementation so that other researchers can replicate our database or refine it for their specific purposes. Our algorithms implemented corpus-based classification, with steps that are easily understandable. Furthermore, this algorithm is efficient in terms of computer processing speed. However, this technique does not take into account the entire context of a tweet and does not account for sarcasm or humor, challenges which still evade the majority of natural language processing algorithms. Our analysis of caloric density of food assumes calories per 100 grams. Most tweets do not specify the exact quantity of food consumed and thus our estimate is just an approximation. Also, caloric expenditure for physical activities was assumed for 30 minutes of physical activity for an individual weighing 155 pounds, which can be an under- or over-estimate depending on the type of activity and persons engaged in that activity.

4.2 Conclusion

There is a growing interest in social determinants of health and how place-based factors affect health. However, one tremendous impediment to such research endeavors includes the lack of open access neighborhood data beyond census demographic and economic characteristics. Moreover, when neighborhood data is available it is inconsistent across geographies and time, thus making neighborhood comparisons difficult, expensive, and time-consuming. Widespread usage of the internet and open recording of many transactions has led to the availability of massive amounts of data that enable understanding of previously hidden micro-level interactions. Social media data uniquely allows us to capture characteristics of the social environment that are understudied. For instance, public posts allow us to potentially measure prevailing sentiment among the people in an area. Additionally public posts about food and physical activity can help

us understand local area social norms and the potential impacts of social modeling of health behaviors.

In this study, we describe algorithms and quality control results for measuring neighborhood characteristics from publicly available, geotagged Twitter data. We demonstrate that tweets can provide a means to assess prevalent sentiment and popularity of types of foods and physical activities in communities, which can enable more efficient targeting of programs and policies to meet the needs of different neighborhoods. In particular, as this study suggests, neighborhoods with social and economic disadvantage as well as those with high percentages of youth may have greater need for health promoting resources.

ACKNOWLEDGEMENTS

This work was supported by NIH grant 5K01ES025433. The funding source did not have any involvement in the study design; collection, analysis and interpretation of data; in the writing of the report; or in the decision to submit the article for publication.

REFERENCES

1. Diez Roux AV. Investigating Neighborhood and Area Effects on Health. *American Journal of Public Health* 2001;**91**(11):1783-1789.
2. Wing S, Barnett E, Casper M, Tyroler HA. Geographic and socioeconomic variation in the onset of decline of coronary heart disease mortality in white women. *American Journal of Public Health* 1992;**82**:204–209.
3. Tyroler HAW, S., Knowles MG. Increasing inequality in coronary heart disease mortality in relation to educational achievement: profile of places of residence, United States, 1962–87. *Annals of Epidemiology* 1993;**3**(Supl):S51–S54.
4. Morris JN, Blane DB, White IR. Levels of mortality, education, and social conditions in the 107 local education authority areas of England. *Journal of Epidemiology and Community Health* 1996;**50**:15–17.
5. Eames M, Ben-Shlomo Y, Marmot MG. Social deprivation and premature mortality: regional comparison across England. *BMJ* 1993;**307**:1097–1102.
6. Townsend P, Phillimore P, Beattie A. *Health and Deprivation. Inequality and the North*. London, England: Routledge, 1988.
7. Waitzman NJ, Smith KR. Phantom of the area: poverty-area residence and mortality in the United States. *American Journal of Public Health* 1998;**88**(6):973-976.
8. Waitzman NJ, Smith KR. Separate but lethal: the effects of economic segregation on mortality in metropolitan America. *Milbank Quarterly* 1998;**76**(3):341-373.
9. Clarke CA, Miller T, Chang ET, Yin D, Cockburn M, Gomez SL. Racial and social class gradients in life expectancy in contemporary California. *Social Science & Medicine* 2010;**70**(9):1373-1380.
10. Truong KD, Ma S. A systematic review of relations between neighborhoods and mental health. *Journal of Mental Health Policy and Economics* 2006;**9**(3):137-154.

11. Wen M, Browning CR, Cagney KA. Poverty, affluence, and income inequality: neighborhood economic structure and its implications for health. *Social Science & Medicine* 2003;**57**:843-860.
12. Mujahid MS, Roux AVD, Shen M, Gowda D, Sánchez B, Shea S, Jacobs DR, Jackson SA. Relation between Neighborhood Environments and Obesity in the Multi-Ethnic Study of Atherosclerosis. *American Journal of Epidemiology* 2008;**167**(11):1349-1357.
13. Black JL, Macinko J, Dixon LB, Fryer JGE. Neighborhoods and obesity in New York City. *Health & Place* 2010;**16**(3):489-499.
14. Heinrich KM, Lee RE, Regan GR, Reese-Smith JY, Howard HH, Haddock CK, Poston WSC, Ahluwalia JS. How Does the Built Environment Relate to Body Mass Index and Obesity Prevalence Among Public Housing Residents? *American Journal of Health Promotion* 2008;**22**(3):187-194.
15. Smith KR, Brown BB, Yamada I, Kowaleski-Jones L, Zick CD, Fan JX. Walkability and body mass index: density, design, and new diversity measures. *American Journal of Preventive Medicine* 2008;**35**(3):237-244.
16. Lysy Z, Booth GL, Shah BR, Austin PC, Luo J, Lipscombe LL. The impact of income on the incidence of diabetes: A population-based study. *Diabetes Research and Clinical Practice* 2013;**99**(3):372-379.
17. Grigsby-Toussaint DS, Lipton R, Chavez N, Handler A, Johnson TP, Kubo J. Neighborhood Socioeconomic Change and Diabetes Risk: Findings from the Chicago Childhood Diabetes Registry. *Diabetes Care* 2010;**33**(5):1065-1068.
18. Morland K, Wing S, Diez RA, Poole C. Neighborhood characteristics associated with the location of food stores and food service places. *American Journal of Preventive Medicine* 2002;**22**:23-29.
19. Stafford M. Pathways to obesity: identifying local, modifiable determinants of physical activity and diet. *Social Science & Medicine* 2007;**65**:1882-97.

20. Wang MC, Kim S, Gonzalez AA, MacLeod KE, Winkleby MA. Socioeconomic and food-related physical characteristics of the neighborhood environment are associated with body mass index. *Journal of Epidemiology Community Health* 2007;**61**(6):491-8.
21. Inagami S, Cohen DA, Finch BK. You are where you shop: grocery store locations, weight, and neighborhoods. *American Journal of Preventive Medicine* 2006;**31**:10–17.
22. Christiansen KMH, Qureshi F, Schaible A, Park S, Gittelsohn J. Environmental Factors That Impact the Eating Behaviors of Low-Income African American Adolescents in Baltimore City. *Journal of Nutrition Education and Behavior* 2013.
23. Morland K, Wing S, Diez Roux A, Poole C. Neighborhood characteristics associated with the location of food stores and food service places. *American Journal of Preventive Medicine* 2002;**22**(1):23-29.
24. Block JP, Scribner RA, DeSalvo KB. Fast food, race/ethnicity, and income: A geographic analysis. *American Journal of Preventive Medicine* 2004;**27**(3):211-217.
25. Roemmich JN, Epstein LH, Raja S, Yin L, Robinson J, Winiewicz D. Association of access to parks and recreational facilities with the physical activity of young children. *Preventive Medicine* 2006;**43**(6):437-441.
26. Brownson RC, Hoehner CM, Day K, Forsyth A, Sallis JF. Measuring the Built Environment for Physical Activity: State of the Science. *American Journal of Preventive Medicine* 2009;**36**(4, Supplement):S99-S123.e12.
27. Mujahid MS, Roux AV, Shen M, Gowda D, Sánchez B, Shea S, Jacobs DJ, Jackson SA. Relation between neighborhood environments and obesity in the multi-ethnic study of atherosclerosis. *American Journal of Epidemiology* 2008;**167**:1349-1357.
28. Oswald AJ, Powdthavee N. Obesity, unhappiness, and the challenge of affluence: theory and evidence. *Economic Journal* 2007;**117**:117:F441–54.
29. Bray I, Gunnell D. Suicide rates, life satisfaction and happiness as markers for population mental health. *Social Psychiatry and Psychiatric Epidemiology* 2006;**41**(5):333-337.

30. Tella RD, MacCulloch RJ, Oswald AJ. The Macroeconomics of Happiness. *The Review of Economics and Statistics* 2003;**85**(4):809-827.
31. Blanchflower DG, Oswald AJ. Hypertension and happiness across nations. *Journal of Health Economics* 2008;**27**(2):218-233.
32. Dodds PS, Harris KD, Kloumann IM, Bliss CA, Danforth CM. Temporal Patterns of Happiness and Information in a Global Social Network: Hedonometrics and Twitter. *PLoS ONE* 2011;**6**(12):e26752.
33. Di Tella R, MacCulloch R. Gross national happiness as an answer to the Easterlin Paradox? *Journal of Development Economics* 2008;**86**(1):22-42.
34. Diez-Roux A. Bringing context back into epidemiology: variables and fallacies in multi-level analysis. *American Journal of Public Health* 1998;**88**:216–22.
35. Macintyre S, Maciver S, Sooman A. Area, class and health: Should we be focusing on places or people. *Journal of Social Policy* 1993;**22**:213–33.
36. Duncan C, Jones K, Moon G. Context, composition and heterogeneity: using multilevel models in health research. *Social Science & Medicine* 1998;**46**:97–117.
37. National Archive of Criminal Justice. Project on Human Development in Chicago Neighborhoods <http://www.icpsr.umich.edu/icpsrweb/PHDCN/> Accessed September 24, 2012.
38. Baltimore Neighborhood Indicators Alliance - The Jacob France Institute. Vital Signs 11 Reports. http://bniajfi.org/vs11_report Accessed September 24, 2013.
39. Peterson RD, Krivo LJ. National Neighborhood Crime Study (NNCS). <http://www.icpsr.umich.edu/icpsrweb/RCMD/studies/27501> Accessed September 25, 2013.
40. Ali MM, Amialchuk A, Heiland FW. Weight-Related Behavior among Adolescents: The Role of Peer Effects. *PLoS ONE* 2011;**6**(6):e21179.
41. Vartanian LR, Sokol N, Herman CP, Polivy J. Social Models Provide a Norm of Appropriate Food Intake for Young Women. *PLoS ONE* 2013;**8**(11):e79268.

42. Cohen DA, Finch BK, Bower A, Sastry N. Collective efficacy and obesity: The potential influence of social factors on health. *Social Science & Medicine* 2006;**62**(3):769-778.
43. Kim D, Subramanian SV, Gortmaker SL, Kawachi I. US state- and county-level social capital in relation to obesity and physical inactivity: A multilevel, multivariable analysis. *Social Science & Medicine* 2006;**63**(4):1045-1059.
44. Berkman L, Syme S. Social networks, host resistance, and mortality: A nine-year follow-up study of Alameda County residents. *American Journal of Epidemiology* 1979;**190**(2):186–204.
45. Bandura A. Social learning theory Englewood Cliffs, New Jersey: Prentice-Hall, 1977.
46. Pachucki MA, Jacques PF, Christakis NA. Social network concordance in food choice among spouses, friends, and siblings. *American Journal of Public Health* 2011;**101**:2170–2177.
47. Keating NL, O’Malley AJ, Murabito JM, Smith KP, Christakis NA. Minimal social network effects evident in cancer screening behavior. *Cancer* 2011;**117**:3045–3052.
48. Rosenquist JN, Murabito J, Fowler JH, Christakis NA. The spread of alcohol consumption behavior in a large social network. *Annals of Internal Medicine* 2010;**152**:426–433.
49. Mednick SC, Christakis NA, Fowler JH. The spread of sleep loss influences drug use in adolescent social networks. *PLoS One* 2010;**e9775**.
50. Roy JP. Socioeconomic status and health: a neurobiological perspective. *Medical Hypotheses* 2004;**62**:222–227.
51. Smith KP, Christakis NA. Social networks and health. *Annual Review of Sociology* 2008;**34**:405-429.
52. Kramer ADI, Guillory JE, Hancock JT. Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences* 2014;**111**(24):8788-8790.
53. Fowler JH, Christakis N. Dynamic spread of happiness in a large social network: Longitudinal analysis over 20 years in the Framingham heart study. *British Medical Journal* 2008;**337**:a2338.

54. Rosenquist JN, Fowler JH, Christakis NA. Social network determinants of depression. *Molecular Psychiatry* 2011;**16**(3):273-281.
55. Bearman PS, Moody J. Suicide and friendships among American adolescents. *American Journal of Public Health*; **94**(1):89–95.
56. Helliwell JF, Layard R, Sachs JD. World Happiness Report 2013. UN Sustainable Development Solutions Network, 2012.
57. Gallup-Healthways. State of American Well-being: 2013 State Rankings and Analysis. <http://info.healthways.com/wbi2013> Accessed January 11, 2014.
58. Quercia D, Ellis J, Capraz L, Crowcroft J. Tracking “Gross Community Happiness” from Tweets. *Association for Computing Machinery Annual Conference on Computer Supported Cooperative Work* Seattle, Washington, 2012; 965-968.
59. Kramer ADI. An unobtrusive behavioral model of "gross national happiness". *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Atlanta, Georgia, USA: ACM, 2010;287-290.
60. Gill AJ, French RM, Gergle D, Oberlander J. The language of emotion in short blog texts. *Proceedings of the 2008 ACM conference on Computer Supported Cooperative Work*. San Diego, CA, USA: ACM, 2008;299-302.
61. Guttman A. R-trees: a dynamic index structure for spatial searching. *Proceedings of the 1984 ACM SIGMOD international conference on Management of Data* 1984;**14**(2):47-57
62. Stanford Natural Language Processing Group. Stanford Tokenizer. <http://nlp.stanford.edu/software/tokenizer.shtml>.
63. Mitchell L, Frank MR, Harris KD, Dodds PS, Danforth CM. The Geography of Happiness: Connecting Twitter Sentiment and Expression, Demographics, and Objective Characteristics of Place. *PLoS ONE* 2013;**8**(5):e64417.

64. Snow R, O'Connor B, Jurafsky D, Ng AY. Cheap and fast-but is it good? Evaluating non-expert annotations for natural language tasks. *Proceedings of the Conference on Empirical Methods in Natural Language Processing* 2008:254-263.
65. United States Department of Agriculture. National Nutrient Database.
<http://ndb.nal.usda.gov/ndb/search/list?format=&count=&max=25&sort=&fg=&man=&lfacet=&qlookup=&offset=50> Accessed February 5, 2014.
66. EnchantedLearning.com. Food and Eating Vocabulary Word List.
<http://www.allaboutsplace.com/wordlist/food.shtml> Accessed February 26, 2014.
67. Zhang N, Campo S, Janz KF, Eckler P, Yang J, Snetselaar LG, Signorini A. Electronic word of mouth on Twitter about physical activity in the United States: exploratory infodemiology study. *Journal of Medical Internet Research* 2013;**15**(11):e261.
68. Ainsworth BE, Haskell WL, Herrmann SD, Meckes N, Bassett DR, Tudor-Locke C, Greer JL, Vezina J, Whitt-Glover MC, Leon AS. 2011 Compendium of Physical Activities: A Second Update of Codes and MET Values. *Medicine & Science in Sports & Exercise* 2011;**43**(8):1575-81.
69. Kendall L, Hartzler A, Klasnja P, Pratt W. Descriptive analysis of physical activity conversations on Twitter. *CHI '11 Extended Abstracts on Human Factors in Computing Systems*. Vancouver, BC, Canada: ACM, 2011;1555-1560.
70. Centers for Disease Control and Prevention. CDC Competition Encourages Use of Social Media to Predict Flu. <http://www.cdc.gov/flu/news/predict-flu-challenge.htm> Accessed February 2, 2014.
71. De Choudhury M, Gamon M, Counts S, Horvitz E. Predicting Depression via Social Media. *Proceedings of the 7th International AAAI Conference on Weblogs and Social Media* Boston, MA, 2013;2.
72. Evans J, Fast S, Markuzon N. Modeling the Social Response to a Disease Outbreak. In: Greenberg A, Kennedy W, Bos N, eds. *Social Computing, Behavioral-Cultural Modeling and*

- Prediction*. Lecture Notes in Computer Science. Vol. 7812 Springer Berlin Heidelberg, 2013;154-163.
73. Lamb A, Paul MJ, Dredze M. Investigating Twitter as a Source for Studying Behavioral Responses to Epidemics. *AAAI Fall Symposium: Information Retrieval and Knowledge Discovery in Biomedical Text*. Arlington, VA, 2012;81-83.
 74. Ghosh D, Guha R. What are we 'tweeting' about obesity? Mapping tweets with Topic Modeling and Geographic Information System. *Cartography and Geographic Information Science* 2013;**40**(2):90-102.
 75. Chen X, Yang X. Does food environment influence food choices? A geographical analysis through "tweets". *Applied Geography* 2014;**51**(0):82-89.
 76. Bliss CA, Kloumann IM, Harris KD, Danforth CM, Dodds PS. Twitter reciprocal reply networks exhibit assortativity with respect to happiness. *Journal of Computational Science* 2012;**3**(5):388-397.
 77. Coulson NS, Buchanan H, Aubeeluck A. Social support in cyberspace: a content analysis of communication within a Huntington's disease online support group. *Patient Education and Counseling* 2007;**68**(2):173-178.
 78. Coulson NS, Knibb RC. Coping with Food Allergy: Exploring the Role of the Online Support Group. *CyberPsychology & Behavior* 2007;**10**(1):145-148.
 79. Verma S, Vieweg S, Corvey WJ, Palen L, Martin JH, Palmer M, Schram A, Anderson KM. Natural Language Processing to the Rescue? Extracting " Situational Awareness" Tweets During Mass Emergency. *Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media*. Barcelona, Spain, 2011;385-392.
 80. Pak A, Paroubek P. Twitter as a Corpus for Sentiment Analysis and Opinion Mining. *Proceedings of the Seventh conference on International Language Resources and Evaluation* Vol. 10, 2010;1320-1326.

81. Widener MJ, Li W. Using geolocated Twitter data to monitor the prevalence of healthy and unhealthy food references across the US. *Applied Geography* 2014;**54**:189-197.
82. Duggan M, Ellison NB, Lampe C, Lenhart A, Madden M. Social Media Update 2014. In: Pew Research Center, ed. Available at: http://www.pewinternet.org/files/2015/01/PI_SocialMediaUpdate20144.pdf, 2015.
83. Burton SH, Tanner KW, Giraud-Carrier CG, West JH, Barnes MD. " Right time, right place" health communication on Twitter: value and accuracy of location information. *Journal of Medical Internet Research* 2012;**14**(6).
84. Morstatter F, Pfeffer J, Liu H, Carley KM. Is the Sample Good Enough? Comparing Data from Twitter's Streaming API with Twitter's Firehose. *arXiv:1306.5204v1 [cs.SI]* 2013.
85. Luo F, Cao G, Mulligan K, Li X. Explore Spatiotemporal and Demographic Characteristics of Human Mobility via Twitter: A Case Study of Chicago. *Applied Geography* 2015;**70**:11-25.
86. Keels M, Duncan GJ, Deluca S, Mendenhall R, Rosenbaum J. Fifteen years later: Can residential mobility programs provide a long-term escape from neighborhood segregation, crime, and poverty. *Demography*; **42**(1):51-73.
87. Cutler DM, Glaeser EL, Vigdor JL. When are ghettos bad? Lessons from immigrant segregation in the United States. *Journal of Urban Economics* 2008;**63**(3):759-774.

Figure Legends

Figure 1. Geography of happiness scores across census tracts in Salt Lake County

Figure 2. Geography of happiness scores across census tracts in San Francisco County

Figure 3. Geography of happiness scores across census tracts in New York County

Table 1. Descriptive statistics of tract-level characteristics

	San Francisco County N=196)	Salt Lake County (n=212)	New York County (n=288)
	Mean (Standard Deviation)	Mean (Standard Deviation)	Mean (Standard Deviation)
Happiness scores (continuous)	6.4 (0.2)	6.2 (0.2)	6.5 (0.2)
<i>Categorized happiness scores</i>			
Percent happy	46.2 (9.0)	37.9 (6.4)	52.3 (8.6)
Percent neutral	43.4 (6.0)	48.1 (4.3)	39.0 (5.8)
Percent sad	10.4 (4.8)	14.0 (3.7)	8.8 (5.0)
<i>Food culture</i>			
Percent of tweets that have food references	6.6 (4.1)	3.1 (1.8)	6.6 (3.8)
Percent of tweets that have healthy food references	1.1 (0.8)	0.5 (0.4)	1.1 (0.7)
Percent of tweets that have fast food restaurant references	0.4 (0.8)	0.3 (0.4)	0.6 (1.0)
Calories density of food mentions (calories per 100 grams)	260.3 (71.3)	261.2 (48.9)	249.6 (56.0)
Sentiment of healthy foods	6.6 (0.4)	6.5 (0.5)	6.7 (0.4)
Sentiment of fast food	6.3 (0.8)	6.4 (0.7)	6.5 (0.5)
<i>Physical activity culture</i>			
Percent of tweets with physical activity references	1.7 (1.4)	1.7 (1.3)	1.7 (1.1)
Caloric expenditure of physical activity references (calories per 30 minutes)	218.5 (47.2)	216.8 (38.5)	218.6 (46.0)
Physical activity is gym-related	17.0 (17.1)	13.3 (12.7)	23.5 (16.4)
Sentiment of physical activity	6.5 (0.4)	6.4 (0.3)	6.6 (0.3)
Census 2010 and ACS 2013 Data			
<i>Population characteristics</i>			
Percent 65 years+	13.8 (7.0)	9.3 (4.8)	13.1 (6.8)
Percent 10-24 years	14.8 (7.5)	22.5 (3.7)	17.6 (9.0)
Percent male	51.0 (5.7)	50.4 (3.5)	47.9 (5.7)
Percent African American	6.6 (9.6)	1.7 (1.5)	17.1 (21.5)
Percent white	50.2 (22.2)	81.5 (12.1)	56.9 (27.0)
Percent Hispanic	14.8 (11.8)	16.5 (13.7)	23.4 (23.1)

Percent relatives (besides spouse and children) living in households	8.0 (6.5)	7.0 (3.3)	5.7 (5.3)
Percent unmarried partner living in households	4.2 (2.5)	2.1 (1.2)	3.5 (1.7)
<i>Household characteristics</i>			
Household size	2.4 (0.7)	3.0 (0.7)	2.0 (0.5)
Percent single female-headed households	9.5 (7.6)	5.9 (2.5)	13.0 (12.1)
Percent householder living alone	36.1 (15.9)	21.7 (13.0)	44.9 (12.6)
<i>Economic characteristics</i>			
Percent owner-occupied housing	37.7 (23.0)	67.7 (21.9)	21.7 (18.4)
Median family income	98893.2 (49650.0)	64267.6 (24449.0)	117429.9 (76118.1)
Percent college graduates	52.1 (21.5)	20.6 (9.5)	58.3 (25.7)
Unemployment rate	8.7 (4.6)	7.7 (3.6)	9.4 (8.3)
Percent less than high school graduate	8.6 (12.1)	15.2 (10.9)	11.4 (14.3)
Percent families living in poverty	8.6 (8.3)	9.9 (8.2)	13.0 (12.9)

All variables are standardized to have a mean of 0 and standard deviation of 1.

Table 2. Census tract demographic and economic predictors of happiness scores

Census tract characteristics	San Francisco County	Salt Lake County	New York County
<i>Population characteristics</i>	Beta (95% CI)	Beta (95% CI)	Beta (95% CI)
Percent 65 years+	0.00 (-0.12, 0.13)	0.18 (0.01, 0.35)*	-0.09 (-0.22, 0.03)
Percent 10-24 years	-0.08 (-0.23, 0.06)	0.02 (-0.13, 0.17)	-0.22 (-0.35, -0.09)*
Percent male	0.03 (-0.11, 0.17)	0.26 (0.05, 0.48)*	-0.04 (-0.21, 0.13)
Percent African American	-0.58 (-0.88, -0.28)*	0.17 (-0.04, 0.38)	-0.07 (-0.21, 0.07)
Percent Hispanic	-0.01 (-0.26, 0.23)	-0.24 (-0.49, 0.01)#	-0.11 (-0.30, 0.08)
Household size	-0.44 (-0.57, -0.30)*	-0.26 (-0.47, -0.05)*	-0.03 (-0.33, 0.27)
Economic disadvantage ^a	0.10 (-0.11, 0.32)	-0.20 (-0.45, 0.04)#	-0.20 (-0.50, 0.09)
<i>Adjusted r-square</i>	0.37	0.28	0.23

N=196 census tracts in San Francisco County, N=212 tracts in Salt Lake County, and N=288 tracts in New York County

All variables are standardized to have a mean of 0 and standard deviation of 1

^aEconomic disadvantage factor score derived from the following census tract characteristics: percent female-headed households, percent families living in poverty, unemployment rate, and median family income (reverse coded)

Adjusted linear regression models includes all predictors simultaneously, and were run separately for each county

*p<0.05 #p<0.10

Table 3. Most popular food items, in descending order of popularity

San Francisco County	Salt Lake County	New York County
coffee	coffee	coffee
beer	beer	starbucks
starbucks	pizza	beer
wine	starbucks	pizza
tea	sushi	wine
pizza	chicken	ice cream
ice cream	ice cream	burger
sushi	chocolate	cheese
chocolate	cake	bbq
burger	cookies	chicken
chicken	wine	tea
cheese	cheesecake	chocolate
bacon	burger	sushi
crab	tea	salad
bbq	tacos	cake
salad	salad	lobster
ramen	eggs	milk
shrimp	fries	ramen
burrito	whiskey	steak
milk	bbq	tacos
tacos	cheese	bacon
steak	burrito	chipotle
cake	taco bell	egg
egg	donut	shrimp
pork	pie	
	milk	
	chips	
	cheese	
	egg	
	mcdonalds	
	pancakes	
	rice	
	donuts	

These food items compose 50% of all food references

Table 4. Census tract predictors of percent of tweets with healthy food references

Census tract characteristics	San Francisco County	Salt Lake County	New York County
<i>Population characteristics</i>	Beta (95% CI)	Beta (95% CI)	Beta (95% CI)
Percent 65 years+	-0.24 (-0.39, -0.09)*	-0.15 (-0.35, 0.05)	-0.13 (-0.28, 0.01)#
Percent 10-24 years	-0.24 (-0.42, -0.07)*	-0.05 (-0.22, 0.12)	-0.15 (-0.30, 0.00)*
Percent male	-0.20 (-0.37, -0.03)*	-0.06 (-0.30, 0.19)	-0.08 (-0.27, 0.11)
Percent African American	-0.44 (-0.80, -0.08)*	-0.06 (-0.31, 0.18)	-0.01 (-0.17, 0.14)
Percent Hispanic	0.22 (-0.07, 0.51)	-0.01 (-0.29, 0.28)	-0.16 (-0.38, 0.06)
Household size	-0.25 (-0.41, -0.09)*	-0.35 (-0.59, -0.11)*	-0.10 (-0.44, 0.24)
Economic disadvantage ^a	0.14 (-0.12, 0.40)	0.00 (-0.28, 0.28)	-0.01 (-0.35, 0.33)
<i>Adjusted r-square</i>	0.14	0.05	0.08

N=196 census tracts in San Francisco County, N=212 tracts in Salt Lake County, and N=288 tracts in New York County

All variables are standardized to have a mean of 0 and standard deviation of 1

^aEconomic disadvantage factor score derived from the following census tract characteristics: percent female-headed households, percent families living in poverty, unemployment rate, and median family income (reverse coded)

Adjusted linear regression models includes all predictors simultaneously, and were run separately for each county

*p<0.05 #p<0.10

Table 5. Most popular physical activity terms, in descending order of popularity

San Francisco County	Salt Lake County	New York County
walk/walking	walk/walking	walk/walking
dance/dancing	dance/dancing	dance
running	hike/hiking	running
yoga	running	yoga
hike/hiking	work out/workout	workout
workout	golf	(swimming) pool
(swimming) pool	skiing	dancing
golf	(swimming) pool	golf
swimming	yoga	
	swim	
	bowling	

These terms compose 75% of all physical activity references

Table 6. Census tract predictors of percent of tweets with physical activity references

Census tract characteristics	San Francisco County	Salt Lake County	New York County
<i>Population characteristics</i>	Beta (95% CI)	Beta (95% CI)	Beta (95% CI)
Percent 65 years+	0.25 (0.10, 0.41)*	-0.05 (-0.24, 0.14)	-0.01 (-0.16, 0.14)
Percent 10-24 years	-0.02 (-0.21, 0.16)	-0.06 (-0.22, 0.11)	-0.15 (-0.30, 0.01)#
Percent male	0.07 (-0.10, 0.25)	0.06 (-0.18, 0.30)	-0.12 (-0.32, 0.07)
Percent black	-0.05 (-0.43, 0.32)	0.07 (-0.17, 0.31)	0.05 (-0.11, 0.21)
Percent Hispanic	-0.06 (-0.37, 0.25)	0.08 (-0.20, 0.36)	0.01 (-0.22, 0.23)
Household size	-0.03 (-0.20, 0.14)	-0.20 (-0.44, 0.04)#	0.37 (0.01, 0.72)*
Economic disadvantage ^a	-0.03 (-0.29, 0.24)	-0.40 (-0.68, -0.12)*	-0.29 (-0.64, 0.06)#
<i>Adjusted r-square</i>	0.05	0.06	0.02

N=196 census tracts in San Francisco County, N=212 tracts in Salt Lake County, and N=288 tracts in New York County

All variables are standardized to have a mean of 0 and standard deviation of 1

^aEconomic disadvantage factor score derived from the following census tract characteristics: percent female-headed households, percent families living in poverty, unemployment rate, and median family income (reverse coded)

Adjusted linear regression models includes all predictors simultaneously, and were run separately for each county

*p<0.05 #p<0.10

Figure 1. Geography of happiness scores across census tracts in Salt Lake County

Figure 2. Geography of happiness scores across census tracts in San Francisco County

Figure 3. Geography of happiness scores across census tracts in New York County