Chapter 8
Big Data and Wellbeing in the Arab World

Aamna Al-Shehhi, Ian Grey and Justin Thomas

Abstract The rapid and widespread usage of social media platforms, such as Twitter, Facebook and Instagram has given rise to unprecedented amounts of user-generated data. This data contains expressions reflecting users thoughts, opinions and affective states. Systematic explorations of this type of data have begun to yield valuable information about a variety of psychological and cultural variables. To date however, very little of this research has been undertaken in the Arab world. It is important to extend this type of macro-level big data analysis across cultures and languages as each situation is likely to present different methodological challenges and to reveal findings particular to the sociocultural context. This chapter examines research—much of it our own—exploring subjective wellbeing in the United Arab Emirates (UAE) using data from Twitter and explores the findings from cross-linguistic analysis of happiness (positive–negative affective patterns of language use) and other variables associated with subjective wellbeing in the region. Additionally, we explore temporal patterns of happiness observed in relation to Ramadan and other events of sociopolitical and religio–cultural significance. The UAE focus is discussed with reference to broader trends in data science, sentiment analysis and hedonometry.

8.1 Introduction

The 21st century has been witness to many technological changes and foremost amongst these has been the accelerated digitisation of people's social lives. The thoughts and feelings of millions of people around the globe are expressed daily on

A. Al-Shehhi
Khalifa University of Science and Technology, Abu Dhabi, UAE
e-mail: aamena.shehhi@gmail.com

I. Grey (✉)
School of Social Sciences, Lebanese American University, Beirut, Lebanon
e-mail: Ian.grey@zu.ac.ae

J. Thomas
Zayed University, Abu Dhabi, UAE
e-mail: justin.thomas@zu.ac.ae

L. Lambert and N. Pasha-Zaidi (eds.), *Positive Psychology in the Middle East/North Africa*, https://doi.org/10.1007/978-3-030-13921-6_8

various online platforms and the enormous growth of social media usage witnessed on Facebook, Twitter and Google has led to the accumulation of huge amounts of data (referred to as "Big Data"). These vast datasets have attracted the attention of social scientists as a potential means to provide a greater understanding of various aspects of human behaviour, and a rapidly expanding body of research is now testifying to the power of Big Data to provide insight into a range of psychological questions. However, not only does the use of Big Data facilitate addressing questions but the answers resulting from its use may be as robust, or even more so, than what is currently associated with existing traditional methodologies. For instance, Schwartz et al. (2013a) analysed the language used in tweets from 1300 different US counties and reported that Twitter content was as predictive of the subjective wellbeing of people living in those counties as that measured by representative surveys, over and above established demographic and socio-economic controls like age, gender, ethnicity, income, and education. Other research suggests that social media platforms like Twitter can capture more information about issues such as county-level atherosclerotic heart disease than robust epidemiological models that include socio-economic status, key demographics and health variables like smoking, diabetes and hypertension (Eichstaedt et al., 2015). Finally, other research suggests that Facebook-language-based predictions of personality seem to be about as good as a friend's rating (Park et al., 2014) and that Twitter language can reflect psychological differences with respect to differing political orientations (Sylwester & Purver, 2015).

The present chapter contributes to this new research agenda by showing how Big Data can improve not only our understanding of subjective wellbeing but how it also provides new means to investigate aspects of wellbeing on a scale not possible until now. We first provide a brief overview of why subjective wellbeing has become an important topic in psychological research and how it is conceptualized, and follow by identifying common methods of analysis used in social media and their respective advantages and disadvantages. Finally, we report several of our key studies based on social media data for the United Arab Emirates (UAE), noting the insights and challenges associated with such research in Arab-speaking countries. We argue that the convergence of the rise in interest in aspects of psychological wellbeing and the availability of large cross-linguistic datascapes offers the potential to both present and address important psychological questions for the Arab World.

8.2 The Beginnings and Proof of Concept

A key question from the outset regarding the relevancy of social media to answer psychological questions is; are there sufficient similarities between online digital behaviour and behaviour in the real world? In answering this, many authors point to a seminal study conducted by Google in 2009 (Google, 2009; Olson, Konty, Paladini, Viboud, & Simonsen, 2013), demonstrating that simply by measuring the frequency of online search queries related to the common flu, in addition to topics such as flu symptoms and medications, the spread of flu through the United States could be accurately tracked (Cook, Conrad, Fowlkes, & Mohebbi, 2011). What was particularly important about the study was not only that the use of search terms correlated with the occurrence of flu on the ground but that Google's estimates of future occurrences of the illness were about two weeks ahead of those made by leading medical authorities in the US, and with a very high degree of accuracy between the estimates provided by Google and the data reported by the Centers for Disease Control. This study served as an initial proof of concept suggesting that both the Internet and its data could be used to measure a variety of variables about people that concerned them at a social level. This research was pivotal in demonstrating that when people search for information online, the data revealed by these search patterns can be reflective of the relative importance of their concerns.

8.3 Psychology and Big Data

While the above study points to a possible convergence between online behaviour and aspects of behaviour in the real world, the question that arises is whether wellbeing can be studied as effectively by accessing or monitoring the behaviour of individuals on digital platforms? A traditional limitation of most psychology studies in this domain is that the number of people involved is small, often no more than 50–150 people. While even the largest studies number a few thousand, the ability of these small numbers to generalize to larger populations remains a major criticism of the field. Researchers ask whether the use of Big Data and techniques similar to those employed in the Google flu study can be used to garner insights into the wellbeing of individuals as the potential to access millions of people in contrast to traditional psychological studies relying on self-report is immense. For instance, Facebook alone has approximately one billion users. Thus, rather than rely on generalizations from small samples to entire populations, which have only a limited degree of success, it is possible to get estimates about entire populations or large proportions of given populations owing to their participation in bigger data sets. Moreover, with so many people, a closer and more nuanced look at different subgroups is possible; for example, understanding how women differ from men, young from old, and so forth. However, there is one key additional advantage; social media users tend to behave in more naturalistic ways than is the case with traditional methodologies like surveys, improving ecological validity. The behaviour of people online appears highly similar to natural settings where they talk about things they care about with their friends in what they consider to be private environments. As such, we analyse the language people use on social media, addressed later.

8.4 Subjective Wellbeing

Over the last two decades, there has been a growing body of work that evaluates human wellbeing based on self-reporting by individuals and groups. Generally referred to as measures of subjective wellbeing (SWB), these studies attempt to measure “satisfaction” with their quality of life (Diener, Suh, Lucas, & Smith, 1999). SWB has been defined as the sum of a person’s affective and cognitive needs and though SWB is often assumed to be restricted to happiness, the term in fact covers a wider range of concepts, comprising at least three components: life evaluation (i.e., satisfaction with one’s life or with specific domains), affect (i.e., the balance between positive and negative emotions) and eudemonia (i.e., the extent to which a person has realised their potential and has a meaningful life) (Diener et al., 1999).

The basis of using SWB as a metric is to measure the extent to which human needs are being met and how individuals perceive their quality of life. Because these measures have been typically based on the judgments of survey respondents rather than on more easily quantifiable inputs of money and material goods, there have always been concerns that these “subjective” measures are not factually based and therefore less valid than “objective” measures like GDP. Yet, as noted by several authors, objective measures such as life expectancy, rates of disease, and GDP are only proxies for wellbeing that have been identified through the subjective judgment of decision-makers; “hence the distinction between objective and subjective indicators is somewhat illusory” (Costanza et al., 2007, p. 268). There is also a concern that there are cultural differences that make it difficult to compare the results across different ethnic, gender, age, religion, and other cultural boundaries (Morrison, Jebb, Tay, & Diener, 2018). However, comparisons of reported wellbeing and per capita GDP have shown that beyond a certain income level, happiness does not increase significantly with additional income and economic gains beyond the threshold no longer correlate with increases in wellbeing (Inglehart, 1997). As a result, ways to help identify as accurately as possible how SWB can be meaningfully assessed, in addition to how these metrics can determine the impact of policies on the lives of individuals are being sought. Measuring subjective wellbeing is not without its challenges; most notably, SWB is not always revealed by people’s behaviour or choices and accessing meaningful data typically requires large representative surveys which are time-consuming and financially costly, as well as limited in coverage, time and space. This eliminates the possibility of measuring SWB with the appropriate time-cycle frequency and to posess local relevance for policymakers that is necessary for decision-making.

8.5 Happiness as Policy Agenda

Happiness has become largely interchangeable with subjective wellbeing, which is to say, an experiential state that contains a globally positive affective tone. To date, social scientists have conceptualised and measured happiness in at least two different ways. One of these is what is termed "affect balance", which examines the relative balance between pleasant and unpleasant emotional states, and is thus primarily the sum of how one feels at different moments. The second, life satisfaction, goes beyond momentary feelings to elicit an evaluative assessment of one's life as a whole (Linley, Maltby, Wood, Osborne, & Hurling, 2009).

Spreading across many countries, the "happiness agenda" is not without detractors. Research has shown that having purpose and meaning in life is associated with a number of positive outcomes such as increases in overall wellbeing and life satisfaction, better mental and physical health, increased resiliency, enhanced self-esteem, and reduced likelihood of depression. In contrast, happiness as an affective state is paradoxical in that the more importance people place on being happy, the more they appear to become unhappy and prone to being depressed (Mauss, Tamir, Anderson, & Savino, 2011). For example, individuals who put most emphasis on being happy report less frequent positive emotions, less satisfaction with life, and substantially more depressive symptoms than people who do not have happiness as a priority. Further, individuals that value happiness the most also seem to report less psychological wellbeing. It seems that the very pursuit of happiness seems to thwart its attainment (Mauss et al., 2011).

Another study asked nearly 400 Americans aged between 18 and 78 whether they thought their lives were meaningful and/or happy (Baumeister, Vohs, Aaker, & Garbinsky, 2013). Examining their self-reported attitudes toward meaning in life, happiness, and their interactions with a number of other variables such as stress levels, spending patterns, and having children, the researchers found that while a meaningful life and happy life overlap, they are different. The researchers report that leading a happy life appears linked with what they conceptualise as being a "taker" while leading a meaningful life corresponded with being a "giver". Those people identified as happy obtained joy from receiving benefits from others while people leading meaningful lives obtained joy from giving to others. Happiness without meaning appeared to characterize a relatively shallow, self-absorbed or even selfish life in which things went well, where needs and desires were easily satisfied, and difficult situations were avoided. They concluded that meaning transcends the self while happiness is about giving the self what it wants.

Collectively, these results fit the overall view that having one's needs satisfied, being able to obtain what one wants and needs, and feeling good more often than bad are central to happiness but have relatively little to do with a meaningful life. People are happy when they get what they want. Meaning, on the other hand, is to be found elsewhere. In sum, happiness is important but it is not all there is to wellbeing, nor is it the only thing that ought to be valued. The above suggests that there are critical nuances in the measurement of subjective wellbeing, which are important considerations when approaching data obtained from social media platforms.

8.6 The Language of Social Media

Every day millions of people share their thoughts and feelings through simple text messages written on various social media platforms such as Facebook and Twitter, with such platforms in general creating an unprecedented amount of written language, significant amounts of which are publicly available. Twitter users alone write approximately 500 million messages every day (Krikorian, 2013). An immediate implication of the corpus of written language accumulating on such platforms is that it represents a massive source of rich psychological data with substantial scientific potential. The task confronting researchers is to develop methods that can meaningfully and effectively access this data. Whereas in the past researchers relied on either historical language samples such as literature or prompted participants to write new text, social media provides researchers with the natural language of millions of people with relative ease.

Several characteristics of social media language render it useful for research. First, social media language is typically written in natural social settings such as homes and schools and captures communication among friends, family members and acquaintances. Second, time-consuming prospective studies, whereby researchers track pre-identified variables over a period of time, are less necessary as data can be retroactively accessed for research purposes. Third, social media users disclose information about themselves at high rates; for many users, a frequent topic of discussion is themselves (Naaman, Boase, & Lai, 2010) and as such, social media users typically present their true selves and not mere idealized versions (Mitja et al., 2010). Social media language is a rich source of data; Twitter is particularly popular because it allows researchers to access massive amounts of data quickly, cost-effectively, and without having to obtain any informed consent from users.

Thus, the fundamental unit of analysis of data from social media is the language people use and the underlying premise is that this written language represents a quantifiable signal that can be used to assess a variety of psychological issues as diverse as personality traits (Bogg, 2017; Sap et al., 2014; Schwartz et al., 2013a) to depression (Coppersmith, Dredze, & Harman, 2014). Social media provides unprecedented access to unforced natural language that affords substantial increases in ecological validity (i.e., the extent to which findings can be generalized to real life settings) compared to that provided by traditional survey methodologies.

8.7 Methods to Investigate Language

Language patterns on social media platforms are generally studied using one of two formats: closed vocabulary or open-vocabulary approaches. With relatively few exceptions, psychological studies have used a closed vocabulary or word-counting approach to analysing such language. This method starts with lists of words that are combined into categories (e.g., pronouns) based on theory, followed by a count of the relative frequency of these words within a body of text. As this approach begins with predefined categories of words, it has been described as closed vocabulary. The closed-vocabulary approach scans the text for keywords that are retrieved from a predefined dictionary. Sometimes, these dictionaries contain only a few dozen highly specific words. Most however, comprise thousands. For example, the LabMT dictionary that is used in several studies contains 10,000 words that were rated by Amazon Mechanical Turk on a scale from sad to happy.

The most popular dictionary in SWB research, however, is the Linguistic Inquiry and Word Count (LIWC; Pennebaker, Chung, Ireland, Gonzales, & Booth, 2007), which comprises more than 6000 words assigned to one or more dimensions. The LIWC was developed to capture multiple psychological dimensions (Pennebaker & King, 1999). Words were first compiled from dictionaries, thesauri, existing questionnaires, and other sources, and then independently rated by three or more judges for inclusion in each category (Pennebaker & Francis, 1996; Pennebaker & King, 1999). The 2001 version includes 64 categories within nine main types (affective, social, sensory, biological, and cognitive processes; verbs; relativity, function and miscellaneous words) and has been translated into 12 languages. An updated version (LIWC2015) appeared in April 2016 (www.liwc.net), automatically counting word frequencies for 64 psychologically relevant categories, such as "function words" (e.g., articles, pronouns, conjunctions), "affective processes" (e.g., happy, cried, nervous), and "social processes" (e.g., pal, friend, talk). The dimensions most frequently used by SWB researchers are positive and negative emotion words and dimensions related to specific emotions such as sadness or anger. Though a simple interface, LIWC nonetheless allows text to be turned into numeric values. The program passes text through a "processor" (i.e., tokenizer and word counter) and provides the frequency that a user mentions each category. Frequencies can be adjusted by the total number of words as users differ in the number of words they write, which is also the probability that any random word in a document belongs to a given category. The relative frequencies can then be summarized descriptively, correlated with other variables, or used as predictors in regression analyses. LIWC is easy to use and widely accepted as a validated method to study language patterns; yet, it comprises a limited number of words and does not work well for colloquial language, leading some to question its validity for the analysis of social media language (Luhmann, 2017).

On the other hand, open-vocabulary approaches are not based on an a priori defined list of words but instead use statistical and probabilistic techniques to identify relevant language patterns or topics (i.e., clusters of words frequently used together) in order to find potentially meaningful clusters of words in large samples of natural language (e.g., Grimmer & Stewart, 2013; O'Connor, Bamman, & Smith, 2011; Schwartz et al., 2013b; Yarkoni, 2010). Open-vocabulary methods do not rely on a priori word or category judgments; rather, they extract a comprehensive collection of language features from the text being analysed. In contrast to closed-vocabulary methods, open-vocabulary methods characterize a language sample by the relative use of (a) single, uncategorized words; (b) non-word symbols (e.g., emoticons, punctuation); (c) multiword phrases; and (d) clusters of semantically related words identified through unsupervised methods, or topics. Because these language features are not identified a priori, these methods can accommodate neologisms and unconventional language use. As such, open-vocabulary methods extract more and richer features of language samples.

As an example, Schwartz et al. (2013a) applied an open-vocabulary approach called differential language analysis (DLA) to a large collection of social media messages and identified 2,000 clusters of words or topics. For example, one topic included the words “love”, “sister”, “friend”, “world”, “beautiful”, “precious”, and “sisters”, and a second topic included “government” “freedom”, “rights”, “country”, “political”, and “democracy”. These topics are generated in a data-driven, “bottom-up” manner, as opposed to the theory-driven, “top-down” methods used in closed-vocabulary approaches. Evidence suggests that open-ended approaches outperform closed vocabulary approaches when developing prediction models of language use associated with gender (Schwartz et al., 2013a). In addition, open-vocabulary methods appear to outperform closed-vocabulary methods when predicting the personality of online users in several studies (e.g., Gill, Nowson, & Oberlander, 2009). Open-vocabulary based predictions of SWB also tend to exhibit greater predictive validity than closed-vocabulary-based predictions. In one study, life satisfaction was less strongly correlated with words identified using a closed-vocabulary approach than with topics identified using an open-vocabulary approach (Luhmann, 2017). However, the downside is that open-vocabulary approaches require more sophisticated statistical skills and computational abilities although they are less labour intensive than closed vocabulary approaches in that they do not require the construction of a priori categories.

Despite the relative advantage of open-vocabulary approaches, closed-vocabulary methods remain popular in analyses of social media language (Golbeck, Robles, & Turner, 2011; Holtgraves, 2011; Sumner, Byers, Boochever, & Park, 2012). Within computer science and related fields, researchers have used closed-vocabulary analyses to study how well social media language can predict a user’s personality (Mitja et al., 2010). For example, Golbeck et al. (2011) used a closed vocabulary approach to analyse the language written in the personal profiles and messages of Facebook users, who also completed personality measures. Relative uses of LIWC word categories (e.g., positive emotions, social processes) were then used as predictors in statistical models where the outcomes were self-reports of personality. When applied to a new sample of users, these models predicted users’ personality traits better than chance, with the authors concluding that “users’ Big Five personality traits can be predicted from the public information they share on Facebook” (Golbeck et al., 2011, p. 260).

8.8 Accessing Social Media Data

The social media analytics process involves four distinct steps which are data discovery, collection, preparation and analysis. Accessing social media data begins with raw data which can include comments, shares, likes, mentions, impressions, messages, profiles and hashtag use. Once a particular dataset is selected (e.g., Twitter, or Facebook), the data are typically accessed through an application programming interface (API), which is a format that enables computer programs to communicate with one another. More importantly, an API is also an interface for researchers to collect data from a given social media service. Through small program scripts, researchers can access the API to retrieve, store, and manipulate digital traces left by the users of a service for subsequent empirical analysis. Part of the attraction of API-based research is that the collection, organization, cleaning, preservation and analysis of data can be automated making APIs highly efficient research tools (Lomborg & Bechmann, 2014). Some social media companies such as Twitter provide publically available API's. However, they typically do not provide full and unlimited access to the entire social media dataset and is restricted to the data provided by the company. As such, API research can suffer from a lack of transparency regarding the data output and quality (i.e., the data that the company selects to make available for use with an API). Comparisons between social media platforms are also made difficult by the fact that there are differences between companies in terms of what they give access to.

To date, most API research has been conducted with Twitter as it has been traditionally quite open in terms of access to its database. Facebook API researchers are required to request permission to collect non-public data from participants through a Facebook app, while Twitter makes a random sample of all public tweets available in real time, which can be accessed through the Twitter API (https://dev.twitter.com/streaming/public) and only requires a Twitter account to access. Upon registering, an API key, API secret, access token, and access secret are received by the user. Following this, a blank Twitter application is created, from which data can be selected and retrieved. This can include tweets, user information, entities (metadata and contextual information), and places. As part of the code, a destination for the data is specified, such as a CSV file or database. Twitter places limits on how much information can be requested each hour as a free user (1% random feed per day; alternatively, data based on a specific criterion, such as geographic location can also be requested). Often, meta-data about each social media post is needed, such as the time it was posted, location, and by whom it was posted. Some information (e.g., time of posting) is easy to extract through the API; other information can be inferred from user profiles.

The data accessed through the API form a database with language data (social media posts and their metadata) and associated outcome variables either at the individual, group, or region level. The language data in its raw form is not suited for quantitative analysis as at this stage it remains only a sequence of characters. The next step, tokenization, refers to the process of breaking up posts or sentences into meaningful tokens or words, which may be known dictionary words, misspellings, punctuation, netspeak (e.g., lol, brb), emoticons (e.g., "<3" is a heart, ":)" is a smiling face), and other variations. Sequences of letters are automatically identified with adjustments made to separate punctuation from words. Related to this is that as the API selected for usage with a dataset influences the kinds of research questions that can be meaningfully asked, it is important that researchers have sufficient computational skills or work alongside others who do in order to use social media APIs.

8.8.1 Social Media and Wellbeing

Analysis of social media data can provide a macro-level perspective on various aspects of human behaviour and related psychological variables. As noted, closed-vocabulary approaches, of which there are several examples, are widely used in social media analysis. In psychological research, closed-vocabulary approaches have most commonly been implemented through the LIWC program (Pennebaker et al., 2007). For example, by applying a priori created lexica across thousands of Facebook users and blogs and millions of word instances, extraversion was positively corelated to using more positive emotion words, whereas neuroticism related to using more negative emotion and swear words (Gill et al., 2009; Sumner, Byers, & Shearing, 2011). In over 140 million words from nearly 20,000 blogs, older and male bloggers tended to use more words related to religion, politics, business, and the Internet. Conversely, younger and female bloggers used more personal pronouns, conjunctions, fun, romance, and swear words (Argamon, Koppel, Pennebaker, & Schler, 2007). Across 16,000 Twitter users and two million tweets, Christians used more religious, positive emotion, and social process words, whereas atheists used more negative emotion and insight words (Ritter, Preston, & Hernandez, 2014). In millions of Facebook posts, positive and negative emotion expressions even appear to be related to local weather reports (Coviello et al., 2014).

As mentioned previously, happiness and life satisfaction are correlated but conceptually and functionally distinct (Diener, Diener, Choi, & Oishi, 2018; Diener, Ng, Harter, & Arora, 2010; Diener, Oishi, & Tay, 2018; Lucas, Diener, & Suh, 1996). To date, most approaches using social media have tapped into emotional aspects of wellbeing (i.e., happiness) and these appear to yield valid indicators of emotional wellbeing, as indicated by studies showing that the use of closed vocabulary approaches fluctuate in weekly and diurnal patterns in a manner similar to self-reported emotional wellbeing, and also by studies reporting moderate correlations of these measures with self-reported emotional wellbeing (Luhmann, 2017). For life satisfaction, however, the correlations between self-reported scores and alternative indicators are typically weaker suggesting that individual-level life satisfaction cannot yet be fully predicted reliably from people's digital traces.

Social media research on happiness has been mostly limited to studying positive and negative affect, also called sentiment analysis. Watson, Clark and Tellegen (1988) present a two-dimensional model of human emotion, where one dimension represents the valence of affective states—positive and negative—and a second dimension reflects the level of physiological arousal—high and low. These states are highly dynamic and influenced by internal and external stimuli, such as the weather, social interactions and our subjective interpretations of such external events. These affective states can be expressed in written or spoken language, reflected in English by words such as "happy", "sad", "excited", or "bored". Increasingly, people use social media platforms such as Twitter to express their current status, including direct or indirect references to their affective states. This phenomenon provides researchers with an opportunity to explore affective states as a function of time and place and concerning a specified attitude object (e.g., religion, politics, or smoking).

Sentiment analysis aims to extract and quantify subjective opinions or feelings from the words people use on social media platforms and can be valuable to those interested in monitoring and promoting community or even national wellbeing. Its techniques have been successfully applied across a range of topics including the exploration of public sentiment across time and in response to major events or political occurrences (Diaz, Gamon, Hofman, Kıcıman, & Rothschild, 2016). This interest has led to the development and refinement of sentiment analytic techniques. For instance, a technique known as the Hedonometer and Valence Shift Word Graphs were used to study the affective linguistic trends of song lyrics and blogs posts over time (Dodds & Danforth, 2010); the authors concluded that the happiness (greater expressed positive affect, relative to negative affect) of song lyrics declined from the 1960 to 1990s and then remained stable after 1995. Other studies have attempted to explore temporal patterns in affective states over varying durations (Dodds, Harris, Kloumann, Bliss, & Danforth, 2011). For instance, one study examined affective rhythms across 84 countries using LIWC and developed a lexicon for measuring positive and negative affect from Twitter data (Tausczik & Pennebaker, 2010). They found largely universal diurnal patterns across countries, such that the early part of the day was associated with heightened positive affect, which reached a peak around 6:00 am and then decreased over the duration of the day. This study also found that the weekday morning positivity pattern shifts by two hours during the weekend. Other Twitter studies exploring the temporal dynamics of happiness using the Hedonometer have also observed pronounced weekly patterns; weekends tend to be more positive than early weekdays.

In a business context, Twitter data and sentiment analytics have been employed to measure customer satisfaction within the hospitality industry. One study conducted with hotels in Las Vegas found that the Twitter-derived sentiment correlated well with hotel rankings provided by guests. Exploring the meteorological relationship to mood, another Twitter study used OpinionFinder sentiment lexicon and Profile of Mood States (POMS) to extract sentiment. This data was correlated with data from U.S. meteorological agencies such as temperature, rainfall, snow depth, and wind. High temperatures were associated with fatigue, anger and reduced depression, while snow was associated with increased depression, and precipitation with decreased tiredness (Li, Wang, & Hovy, 2014).

In summary, the sentiment analytic Twitter studies have produced findings that are consistent with intuition and highly convergent with findings obtained via other research methods and data sources. Yet, the majority of this work has focused on the English language. There is a need to extend such inquiry further and explore the validity and reliability of sentiment analytic techniques across sociocultural and linguistic contexts.

8.8.2 Beyond Sentiment Analysis

Beyond transient affective states such as happiness, relatively few studies have focused on self-referential tweets using a template driven retrieval strategy to explore life satisfaction, a more trait-like component of subjective wellbeing. Twitter-derived estimates of life satisfaction appear relatively stable across time (i.e., no temporal patterning) and uninfluenced by seasonal transition, celebrity deaths or political crises (Yang & Srinivasan, 2016). Additionally, Twitter users categorized as "satisfied" as opposed to "dissatisfied" showed patterns in their tweets consistent with previous findings in the subjective wellbeing literature. For example, the "satisfied" group expressed significantly more positive and less negative affective words. They also used less profanity and were significantly more positive about religion than their "dissatisfied" counterparts. The research of Schwartz et al. (2013b) is particularly important when moving beyond sentiment analysis. They analysed tweets originating in 1293 US counties for which results from life satisfaction surveys were independently available. Using lexicon features (LIWC and PERMA, a special lexicon founded on Seligman's (2011) wellbeing research) and LDA derived topic features, they built models to predict life satisfaction at the county level. Consistent with general life satisfaction research in psychology, terms and topics about physical activity and social engagement were positively related to subjective wellbeing, while disengagement words (sleepy, tired) related negatively to wellbeing.

A recent study by Yang and Srinivasan (2016) offers a good example of a top-down or theory-driven approach to examining life satisfaction using a large Twitter two-year dataset. Using the Satisfaction with Life Scale (SWLS) (Diener, Emmons, Larsen, & Griffin, 1985) they derived a set of template-driven, retrieval strategies to obtain tweets conveying self-ratings of life satisfaction. Life satisfaction has received relatively little attention in social media research, largely due to difficulties in operationalizing the term. The SWLS scale (Diener et al., 1985) provided the basis for lexicon retrieval and analysis, and has been cited more than 9800 times and used in many areas, including positive psychology. Starting with a SWLS survey statement (Step 1), the authors first obtained an initial set of equivalent statements through crowdsourcing with MTurk (Step 2). For example, the sentence "my life is peachy" is synonymous with the SWLS statement 3 (I am satisfied with my life). Next, they manually generalized the statements into templates (Step 3) which identified equivalent expressions. Identified words in a template could be substituted by a variety of synonyms from a lexicon or a dictionary (e.g., happy, delighted, content are synonymous in this context). Two sets of templates to retrieve tweets expressing life satisfaction versus life dissatisfaction were developed. Dissatisfaction templates had their own synonym sets, e.g., "sad, depressing, miserable". Then, they built 16 retrieval strategies (search queries) from each template (Step 4). The strategies differed along two dimensions: (1) the number of intervening words allowed in a template (P) and (2) the number of words allowed before or after a text segment that satisfies the template asking questions ("Is my life good?") were filtered out (Yang & Srinivasan, 2016). Using this methodology, the authors were able to extract expressions of satisfaction and dissatisfaction with life. Consistent with their definitions, trends in life satisfaction posts appear resilient to external events (political, seasonal, etc.) unlike affect trends reported by previous researchers (Dodds et al., 2011). Comparing users, they found differences between satisfied and dissatisfied users in several linguistic, psychosocial and other features. For example, the latter post more tweets expressing anger, anxiety, depression, sadness and on death.

Twitter is by no means the only source of big data to help identify relationships between wellbeing and other variables as it also appears that people's use of search terms on Google reveals much about wellbeing. In a key study, Algan and colleagues (2016) built a wellbeing indicator from a combination of search keyword groups derived from Google Trends. Search terms were grouped into twelve domains that reflected three aspects of wellbeing: material conditions (job search, job market, financial security and home finance), social aspects (family stress, family time, civic engagement and personal security), and finally, health and wellness (healthy habits, health conditions, summer activities and education, and ideals). A key methodological strength of the study however was the use of factor analysis to create meaningfully composite terms. For instance, the category family stress contained the following component keywords (domesticabuse, marriagehelp, marriageprob, marriagecounseling, familysupport, custody and womensshelter). The researchers report that a number of search categories were consistently associated with lower subjective wellbeing such as job search, financial security, and family stress. Conversely, search terms related to civic engagement and healthy habits were positively related to wellbeing. Overall, this research suggests that the online use of key search terms can be used to build an index of subjective wellbeing.

8.9 Big Data, Wellbeing and the Gulf Countries

To date, few studies have investigated aspects of wellbeing using data from social media platforms in the Gulf countries. However, given the widespread use of social media in these countries, considerable potential exists to investigate various aspects of wellbeing (Sulaiman & Aquil, 2015). A handful of studies have been conducted suggesting that the use of this data may render meaningful insights. We describe these studies in detail and examine their implications for the future of wellbeing research in the region.

8.9.1 Hedonometer: Using Both Arabic and English

Al-Shehhi et al. (in press) examined the affective component of subjective wellbeing (happiness) in the UAE using a large 2015 dataset extracted from Twitter. More than 8 million Twitter messages (tweets) written in Arabic and English were analysed. Numerous cross-linguistic differences were observed, including significant differ-

The technique is based on using 10222 English words to calculate the overall emotional content of a given text or dataset. The 10222 words had previously been rated for happiness from 1 (least happy) to 9 (most happy). The equation below explains how the total happiness of a given text or tweet is computed.

$$h_{avg}(T) = \frac{\sum_{i=1}^{N} h_{avg}(w_i) f_i}{\sum_{i=1}^{N} f_i} = \sum_{i=1}^{N} h_{avg}(w_i) p_i \quad \ldots\ldots.. \ (1)$$

Where:

T *is a given text*

N *is unique words in* T

w_i *is a given word in the text*

f_i *is the frequency of the ith word* w_i

Δh_{avg} *is the range of words to exclude*; $5 - \Delta h_{avg} < h_{avg} < 5 + \Delta h_{avg}$

$h_{avg}(w_i)$ *average happiness of word* w_i, *it is a natural tuning parameter; by default, we can choose* $\Delta h_{avg}=1$ *to, balance the sensitivity* versus robustness.

$p_i = \frac{f_i}{\sum_{i=1}^{N} f_i}$ is the normalized frequency

Fig. 8.1 The Hedonometer Algorithm based on Dodds et al. (2011)

ences in the number of tweets, followers and friends associated with users of English and Arabic. The main purpose of this study was to explore "happiness", that is, social media sentiment (the use of positive and negative words on social media). More specifically, the study aimed to explore whether the use of such terms would correspond with times when people report feeling more or less positively. This study, the first in the Arabian Gulf region, used the hedonometer 2.0 algorithm detailed in Fig. 8.1.

Using the above-mentioned lexicon-based sentiment analytic tool (Hedonometer) and applying it to English and Arabic, we were able to explore temporal patterns in happiness (sentiment). The findings were rich and corresponded closely to intuitive hypotheses. For example, in Arabic and in English, Friday was the "happiest day" (see Fig. 8.2). In the UAE, Friday is the start of the weekend when schools, colleges and many businesses take a break.

Interestingly, the low point of the week in Arabic was Sunday (the start of the working week in the UAE); in English however (see Fig. 8.2), Monday was the low sentiment point. This relates to the large number of expatriates working in the UAE,

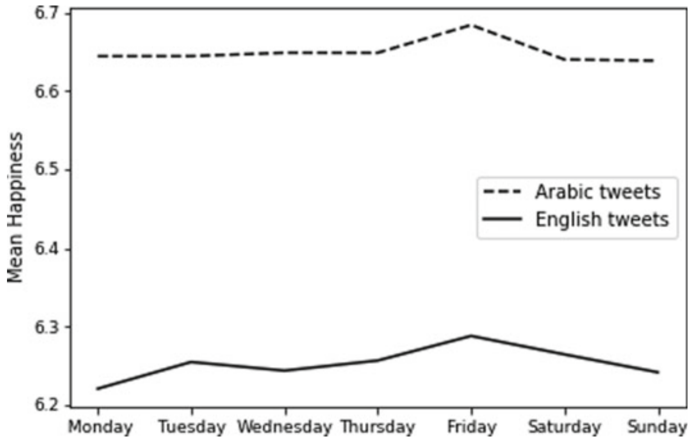

Fig. 8.2 UAE Twitter sentiment (happiness) by days of the week

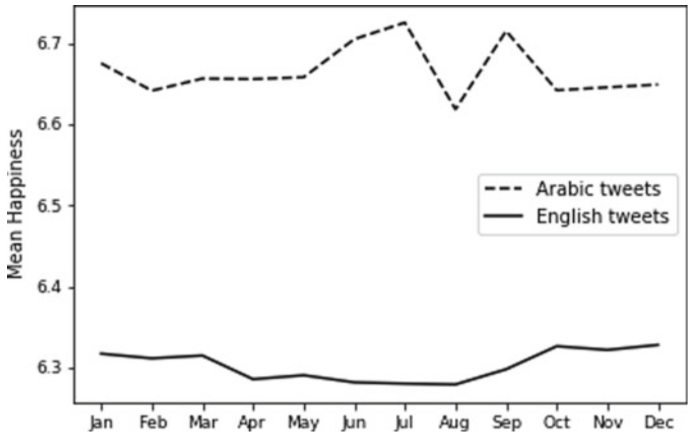

Fig. 8.3 UAE Twitter sentiment (happiness) by months

who come from nations where Monday might be the start of the week given their work with international offices. One interpretation of this finding is that among expatriates (English users), something like an emotional body clock has not yet recalibrated to the new working environment, where the workweek starts on Sunday and not Monday as it did back home. Other findings indicated that 7:00 am was the happiest hour, but may relate to the fact that individuals are tweeting more prior to starting their work day and less as the day progresses. The happiest months differ based on language, and notable patterns are also observed in relation to Ramadan and other events of socio-political and religio-cultural significance in the nation (Fig. 8.3).

The UAE becomes unbearably hot in the summer; people are mostly restricted to indoor activities and this may account for the low points in the English language tweets observed in June through August. It is also when many expatriates leave for their summer and school breaks. In Arabic however, there is a spike in June and July and again in September, with August becoming a low point. This pattern is explained in terms of the Holy month of Ramadan and its corresponding two Eid holidays. In 2015, Ramadan took place during June with Eid al Fitr (a celebration marking the end of Ramadan) in July. In September, there was another big holiday, marking the end of the Hajj pilgrimage (Eid al Adha). These patterns overlap with religio-cultural events being associated with elevated happiness as expressed through social media. The happiest day of the year in the English tweets was New Year's Eve; in Arabic, it was Eid al-Fitr. It is worth noting that religion is still widely practised among the indigenous communities of the Arabian Gulf and is also an important source of national identity, social identity and wellbeing (Thomas, 2014).

8.9.2 Religiosity

Religion often provides people meaning in their lives and is widely viewed as contributing to wellbeing for many individuals (Smith, McCullough, & Poll, 2003). In a study of UAE Twitter data, Thomas, Al-Shehhi and Grey (2018) explored religiosity and the use of profanity among the UAE's Twitter users. Religiosity, similar to happiness, was operationalised as a lexicon, comprising words that unambiguously associated with Islamic religious practice (e.g., Masjid, Mosque, Salah, and Wudu). In this lexicon-driven, cross-linguistic (Arabic/English) analysis of 152 million tweets, the team first looked at differences in language use among users. They found that Arabic users sent a lot more tweets than English users (see Table 8.1).

Although Arabic users represent only 16% of the sample, they account for more than 50% of the tweets. This pattern might be related to the demography of the UAE, where the Arabic speaking citizens are a minority, but also a relatively young minority. UAE citizens comprise around 12% of the population, with a younger median age than the expatriate population (National Bureau of Statistics (NBS), 2009). The discrepancy in tweet frequency by language might also reflect cultural differences (individualism/collectivism) in internet usage (see Würtz, 2005), or

Table 8.1 Tweets by language and unique users

	Tweets	Tweets (%)	Unique users	Users (%)
Arabic	82,065,108	53.69	88,280	16.61
English	42,235,907	27.63	404,689	76.15
Other[a]	28,534,712	18.67	38,444	7.23

[a]There were 64 languages recorded. After Arabic and English, Urdu, Tagalog and Hindi were most frequently used

cultural differences in communication style, high versus low context communication (Hall, 2000), with culture manifesting itself in the online world; however, we are only just beginning to investigate this phenomenon and its implications for wellbeing.

Beyond the linguistic differences in tweet frequency, this study also found patterns of religiosity that support intuitive hypotheses and others that challenge conventional wisdom. Regarding religiosity, Friday (Islam's holy day) was associated with an increase in the overall proportion of religiosity related vocabulary used in Arabic. Similarly, Ramadan (Islam's holiest Month) was associated with a substantial increase in religious sentiment (English and Arabic) and a month-long decrease in the frequency of Arabic obscenity (see Fig. 8.4).

As highlighted in Fig. 8.4, religiosity peaks during the first week of Ramadan then quickly returns to pre-Ramadan levels for much of the month, before again rising during the last week. This pattern seems analogous to the idea of the "hedonic treadmill", where an event (e.g., receipt of a gift) gives rise to elevated positive affect, before a rapid return to a set point (Diener, Lucas, & Scollon, 2006). The dawning of the event, in this case, Ramadan, gives rise to heightened religiosity, which quickly returns to baseline levels (religiosity set-point). The second rise in religiosity during the last week of Ramadan might reflect the religious significance of Laylat Al-Qadr (The Night of Power). The Night of Power is said to be the night when the first verses of the Quran were revealed to the Prophet Muhammad and Muslims are encouraged to perform a special night vigil prayer (Qiam ul Lail) at this time, an activity associated with an abundance of blessings. According to the Quran, "The night of Qadr is better than a thousand months" (97:3). Nobody knows when Laylat Al-Qadr is, however, it is agreed that it falls on one of the last ten nights of Ramadan. The seeking of Laylat Al-Qadr explains the resurgence of religiosity in the UAE Twitter data corresponding with the last week of Ramadan.

This study also explored gender differences, based on a relatively small subset of the data (confirmed males and females), finding some expected and unexpected

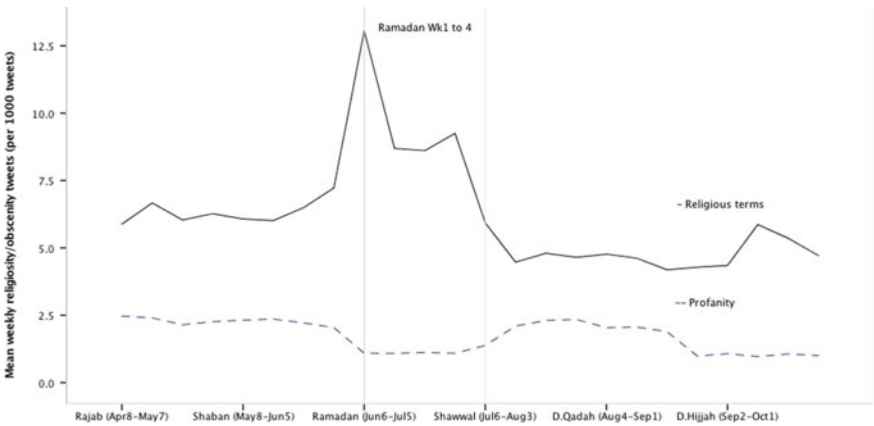

Fig. 8.4 UAE Twitter sentiment (Arabic religiosity) by months

gender differences. Gender was assumed by Twitter display names, i.e., Fatima/Mariam versus Mohammed. Proportionally, females (Fatima's and Mariam's) tweeted in English more than males (Mohammed's). This might be explained in terms of females outperforming males in tertiary education in the UAE. Almost all colleges and universities in the UAE use English as the language of tuition and previous research has documented greater female participation in tertiary education, as well as higher academic performance (Thomas, Al-Marzooqi, & Raynor, 2012). It might be that female Twitter users in the UAE have greater English language proficiency and feel more comfortable tweeting in English. However, this linguistic gender difference could equally be framed as males being more proficient in Arabic and perhaps less Western acculturated. At least one previous study in the UAE has suggested that Western acculturation is greater among UAE females (Ghubash, Daradkeh, Al-Muzafari, El-Manssori, & Abou-Saleh, 2001).

In addition to tweeting more in Arabic, males also expressed more religiosity, both in English and Arabic. This greater public expressivity of religiosity may reflect traditional gender roles in Islam, where males are encouraged to pray in the mosque and women encouraged to pray at home. Similarly, Islamic public service roles, such as Imam (prayer leader), have typically been restricted to men. That said, past research typically reports elevated religiosity among women (Miller & Stark, 2002; Schnabel, 2015). Although previous research has focused on Christian denominations, some studies show a similar, although less pronounced, pattern among Muslims (Stark, 2002). The present atypical findings (females expressing less religiosity than males) are interpretable within the socialisation hypothesis of female religiosity, which suggests that social expectations contribute to females reporting greater religiosity than males (Schnabel, 2015). Social media, however, affords people a degree of anonymity and this might lessen the pressure to conform to social expectations. There are alternative interpretations. For example, among Muslims, female social media users might be highly atypical of the general Muslim female population of the UAE, while their male counterparts are more representative. Such a situation might arise due to pervasive gender-role expectations concerning social interaction. For example, within some households, it might be argued that a "good Muslim woman" does not speak to, or share information with strangers on the Internet.

This same study also examined the use of profanity (obscenities). Although often used casually, profanity can be associated with the expression of negative affect (anger). Looking at the Twitter data in Arabic, revealed patterns are consistent with previous research exploring English curse-words (Wang, Chen, Thirunarayan, & Sheth, 2014), where males used more obscenities than females. This pattern, however, was reversed in English, where females proportionally used more obscenity than males. In cases where English was a second language, using English obscenity might be less emotive. Previous research among bilingual (Turkish/English) speakers has found that using obscenity (e.g., whore) in the second language evoked significantly less autonomic reactivity (emotion) than it did in the first language (Harris, Aycicegi, & Gleason, 2003). This affective attenuation (dampening of emotion) associated with second language use might also explain the far higher rate of English obscenity across the whole dataset. The UAE is a bilingual nation with English widely spoken in education and at the workplace, while Arabic remains widely spoken at home by UAE citizens and Arab expatriates (Mourtada-Sabbah, Al-Mutawa, Fox, & Walters, 2008). This study is important in highlighting some of the variables that need to be considered when exploring wellbeing in the context of a Gulf nation.

8.9.3 I Versus We: Exploring Self-concept on Twitter Using Arabic and English

Past research has indicated that a strong sense of belonging and the expectation of social support act as a buffer to self-esteem and wellbeing (Haslam, Jetten, Postmes, & Haslam, 2009). Viewing oneself (self-concept) as belonging to valuable social groups is another important variable potentially impacting wellbeing. The cultural influence hypothesis suggests that people from individualistic societies tend to emphasise more personal attributes (tall, athletic, geeky) when articulating their self-concepts, while individuals from collectivist societies place greater emphasis on social attributes (Emirati, Muslim, Father). A study by our team (Thomas et al., in prep) explored this hypothesis based on an extract of Twitter data for the UAE. In the first study, the team randomly extracted a thousand user biographies. A user biography (bio) is a short narrative statement, typically written by the user, about the user, which can be viewed as a snapshot of the individual's self-concept. Bios were explored using structured content analysis and the differences in self-concept between English and Arabic users were quantified. In a second study, using an expanded dataset, the team used the LIWC software to explore differences between individuals with biographies uniquely characterised by either group (we, our, us) or self (I, mine, me) referential linguistic constructs ($N = 62{,}038$). In line with the study's predictions, Arabic users mentioned more social attributes than their English speaking counterparts. Conversely, English speakers mentioned more personal attributes. Study Two found that self-referentiality was associated with the more frequent mention of personal attributes, while group-referentiality (us, our, we) was associated with social attributes. Participants who were more social in their orientation (bios mentioning we, us, our) also had tweets that expressed more positive and less negative affect than their self-referential counterparts. In sum, the Big Data extracted from social media platforms such as Twitter, offers a potentially useful means of exploring wellbeing relevant constructs across cultures.

8.10 Challenges to Big Data Analysis in Arab Countries

While the use of Big Data to formulate and answer psychological questions, particularly around aspects of wellbeing in Arabic speaking countries is in its infancy, there are two challenges confronting researchers in the Arab world. Firstly, there is an absence of dictionaries similar to the LIWC for the easy analysis of Arabic. At present, researchers must develop dictionaries for closed language analytic methods by hand. The second challenge is that many of the words used are simultaneously commonly used forenames (polysemy). For instance, Sa'eed is the Arabic word for happy yet is also a common male forename. As a result, using words such as happy in sentiment analysis is made highly problematic resulting in many false positives. Psycholinguists and data scientists will need to collaborate on developing more sophisticated search and quantification strategies for the Arabic language for progress to be made.

8.11 Conclusion

Big Data has great potential in the study of wellbeing. Traditional survey methodologies for the assessment of wellbeing are beset by a range of problems relating to response biases and heuristics such as acquiescence, extreme responding, random responding, digit preferences, primacy and recency effects, demand characteristics, consistency bias and priming effects (Miller, 2017). Though not without its methodological considerations, Big Data at its core relies on natural language and is not particularly prone to these considerations. Social media analysis, such as those reported here, might help further explore the rate and direction of social change in rapidly developing Arabian Gulf states such as the UAE. These surveillance methods offer obvious advantages over self-report survey methodologies in that the data are less influenced by socially desirable responding and other forms of reactivity (Ritter et al., 2014). Social media use frequently involves individuals conversing with one another on wide-ranging topics, freely expressing attitudes, fears and aspirations. In addition to this relative ecological validity, the sample sizes that can be obtained through social media are unparalleled in direct survey work. It is this sample size that silences many potential criticisms, such as the fact that some tweets are posted by news agencies, or commercial entities. Similarly, some tweets are not a reflection of a person's thoughts or emotional state as they might simply be a quoting a celebrity or the lyrics of a song, etc. Big Data however, will typically ensure that such noise is overwhelmingly drowned out by the mass of data that is reflective of people's genuine concerns and sentiment (Dodds et al., 2011). When it comes to helping social scientists learn more about happiness and wellbeing in the MENA/GCC region, the future of Big Data is bright indeed.

References

Algan, Y., Beasley, E., Guyot, F., Higa, K., Murtin, F., & Senik, C. (2016, May 10). *Big data measures of wellbeing: Evidence from a Google wellbeing index in the United States.* OECD Statistics Working Papers, 2016/03. Paris, France: OECD Publishing.

Al Shehhi, A., Thomas, J., Welsch, R., Grey, I. M., Woon, W. L., & Aung, A. (in press). Arabia Felix 2.0: A cross-linguistic twitter analysis of the temporal patterns of happiness in the United Arab Emirates. *EPJ Data Science Journal*.

Argamon, S., Koppel, M., Pennebaker, J., & Schler, J. (2007, September). Mining the blogosphere: Age, gender and the varieties of self-expression. *First Monday*, *12*(9).

Baumeister, R., Vohs, K., Aaker, J., & Garbinsky, E. (2013). Some key differences between a happy life and a meaningful life. *The Journal of Positive Psychology, 8*(6), 505–516.

Bogg, T. (2017). Social media membership, browsing, and profile updating in a representative U.S. sample: Independent and interdependent effects of big five traits and aging and social factors. *Frontiers in Psychology*, *8*, 1122.

Cook, S., Conrad, C., Fowlkes, A., & Mohebbi, M. (2011). Assessing Google flu trends performance in the United States during the 2009 influenza virus A (H1N1) pandemic. *PLoS ONE, 6*(8), e23610.

Coppersmith, G., Dredze, M., & Harman, C. (2014). Quantifying mental health signals in Twitter. In *Proceedings of the Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality* (pp. 51–60).

Costanza, R., Fisher, B., Ali, S., Beer, C., Bond, L., Boumans, R., … Snapp, R. (2007). Quality of life: An approach integrating opportunities, human needs, and subjective wellbeing. *Ecological Economics*, *61*(2–3), 267–276.

Coviello, L., Sohn, Y., Kramer, A., Marlow, C., Franceschetti, M., & Christakis, N. (2014). Detecting emotional contagion in massive social networks. *PLoS ONE, 9*(3), e90315.

Diaz, F., Gamon, M., Hofman, J., Kıcıman, E., & Rothschild, D. (2016). Online and social media: Data as an imperfect continuous panel survey. *PLoS ONE, 11*(1), e0145406.

Diener, E., Suh, E. M., Lucas, R. E., & Smith, H. L. (1999). Subjective wellbeing: 3 decades of progress. *Psychological Bulletin, 125*, 276–302.

Diener, E., Diener, C., Choi, H., & Oishi, S. (2018a). Revisiting "Most People Are Happy"—And discovering when they are not. *Perspectives on Psychological Science, 13*(2), 166–170.

Diener, E., Emmons, R. A., Larsen, R. J., & Griffin, S. (1985). The satisfaction with life scale. *Journal of Personality Assessment, 49,* 71–75.

Diener, E., Oishi, S., & Tay, L. (2018b). Advances in subjective wellbeing research. *Nature Human Behavior, 2*(4), 253–260.

Diener, E., Lucas, R. E., & Scollon, C. N. (2006). Beyond the hedonic treadmill: Revising the adaptation theory of wellbeing. *American Psychologist, 61*(4), 305–314.

Diener, E., Ng, W., Harter, J., & Arora, R. (2010). Wealth and happiness across the world: Material prosperity predicts life evaluation, whereas psychosocial prosperity predicts positive feeling. *Journal of Personality and Social Psychology, 99*(1), 52–61.

Dodds, P., & Danforth, C. (2010). Measuring the happiness of large-scale written expression: Songs, blogs, and presidents. *Journal of Happiness Studies, 11*(4), 441–456.

Dodds, P. S., Harris, K. D., Kloumann, I. M., Bliss, C. A., & Danforth, C. M. (2011). Temporal patterns of happiness and information in a global social network: Hedonometrics and Twitter. *PLoS ONE, 6*(12), e26752, 1–26.

Eichstaedt, J. C., Schwartz, H. A., Kern, M. L., Dziurzynski, L., Lucas, R. E., Agrawal, M., … Ungar, L. H. (2013). Characterizing geographic variation in wellbeing using tweets. In *Proceedings of the 7th International AAAI Conference on Weblogs and Social Media (ICWSM), July 8–1, 2013*, Boston, MA.

Eichstaedt, J. C., Schwartz, H. A., Kern, M. L., Park, G., Labarthe, D. R., Merchant, R. M., … Seligman, M. E. P. (2015). Psychological language on Twitter predicts county-level heart disease mortality. *Psychological Science*, *26*(2), 159–169.

Gill, A. J., Nowson, S., & Oberlander, J. (2009). What are they blogging about? Personality, topic and motivation in blogs. In E. Adar, M. Hurst, T. Finin, N. S. Glance, N. Nicolov, & B. L. Tseng (Eds.), *Proceedings of the 3rd International ICWSM Conference* (pp. 18–25). The AAAI Press.
Golbeck, J., Robles, C., & Turner, K. (2011). Predicting personality with social media. In *Proceedings of the 2011 Annual Conference on Human Factors in Computing Systems* (CHI EA 2011; May 7–12, 2011) (pp. 253–262). Vancouver, BC, Canada.
Google. (2009). *Google flu trends*. Retrieved from www.google.org/flutrends.
Grimmer, J., & Stewart, B. (2013). Text as data: The promise and pitfalls of automatic content analysis methods for political texts. *Political Analysis, 21*(3), 267–297.
Ghubash, R., Daradkeh, T. K., Al-Muzafari, S. M. A., El-Manssori, M. E., & Abou-Saleh, M. T. (2001). Al Ain community psychiatric survey IV: Socio-cultural changes (traditionality-liberalism) and prevalence of psychiatric disorders. *Social Psychiatry and Psychiatric Epidemiology, 36,* 565–570.
Hall, E. T. (2000). Context and meaning. In L. Samovar & R. E. Porter (Eds.), *Intercultural communication: A reader* (pp. 34–43). Belmont, CA: Wadsworth Publishing Co.
Harris, C. L., Aycicegi, A., & Gleason, J. B. (2003). Taboo words and reprimands elicit greater autonomic reactivity in a first language than in a second language. *Applied Psycholinguistics, 24,* 561–579.
Haslam, S. A., Jetten, J., Postmes, T., & Haslam, C. (2009). Social identity, health and wellbeing: An emerging agenda for applied psychology. *Applied Psychology: An International Review, 58*(1), 1–23.
Holtgraves, T. (2011). Text messaging, personality, and the social context. *Journal of Research in Personality, 45*(1), 92–99.
Inglehart, R. (1997). *Modernization and postmodernisation: Cultural, economic and political change*. Boston, MA: Princeton University Press.
Krikorian, R. (2013, August 16). *New Tweets per second record, and how!* (Web Log Post). Retrieved May 14, 2018, from https://blog.twitter.com/engineering/en_us/a/2013/new-tweets-per-second-record-and-how.html
Li, J., Wang, X., & Hovy, E. (2014). What a nasty day: Exploring mood-weather relationship from Twitter. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management, CIKM 2014* (pp. 1309–1318). New York, USA: ACM.
Linley, P., Maltby, J., Wood, A., Osborne, G., & Hurling, R. (2009). Measuring happiness: The higher order factor structure of subjective and psychological wellbeing measures. *Personality and Individual Differences, 47*(8), 878–884.
Lomborg, S., & Bechmann, A. (2014). Using APIs for data collection on social media. *The Information Society, 30*(4), 256–265.
Lucas, R. E., Diener, E., & Suh, E. (1996). Discriminant validity of wellbeing measures. *Journal of Personality and Social Psychology, 71*(3), 616–628.
Luhmann, M. (2017). Using Big Data to study subjective wellbeing. *Current Opinion in Behavioral Sciences, 18,* 28–33.
Mauss, I. B., Tamir, M., Anderson, C. L., & Savino, N. S. (2011). Can seeking happiness make people happy? *Paradoxical effects of valuing happiness. Emotion, 11*(4), 807–815.
Miller, A., & Stark, R. (2002). Gender and religiousness: Can socialization explanations be saved? *American Journal of Sociology, 107*(6), 1399–1423.
Miller, P. (2017). Is there a future for surveys? *Public Opinion Quarterly, 81,* 205–212.
Mitja, D., Back, J., Vazire, S., Gaddis, S., Schmukle, S., Egloff, B., et al. (2010). Facebook profiles reflect actual personality, not self-idealization. *Psychological Science, 21*(3), 372–374.
Morrison, M., Jebb, A., Tay, L., & Diener, E. (2018). Subjective wellbeing across the lifespan worldwide: The role of culture and changing priorities. Under review with *Psychology and Aging.*
Mourtada-Sabbah, N., Al-Mutawa, M., Fox, J. W., & Walters, T. (2008). Media as social matrix in the United Arab Emirates. In A. Alsharekh & R. Sprinborg (Eds.), *Popular culture and political identity in the Arab Gulf States*. London, UK: SAQI in association with The London Middle East Institute at SOAS.

Naaman, M., Boase, J., & Lai, C. (2010). Is it really about me? Message content in social awareness streams. *Proceedings of the ACM Conference on Computer Supported Cooperative Work, CSCW* (pp. 189–192).

NBS. (2009). *Population estimates (2006–2009)* (p. 10). United Arab Emirates: UAE National Bureau of Statistics.

O'Connor, B., Bamman, D., & Smith, N. A. (2011). Computational text analysis for social science: Model assumptions and complexity. In *2nd Workshop on Computational Social Science and Wisdom of the Crowds (NIPS, 2011)*. Retrieved from http://repository.cmu.edu/cgi/viewcontent.cgi?article=1201&context=lti.

Olson, D. R., Konty, K. J., Paladini, M., Viboud, C., & Simonsen, L. (2013). Reassessing Google flu trends data for detection of seasonal and pandemic influenza: A comparative epidemiological study at three geographic scales. *PLoS Computational Biology, 9*(10), e1003256.

Park, G., Schwartz, H. A., Eichstaedt, J. C., Kern, M. L., Kosinski, M., Stillwell, D. J., … Seligman, M. E. P. (2014). Automatic personality assessment through social media language. *Journal of Personality and Social Psychology*, *108*(6), 934–952.

Pennebaker, J., Chung, C., Ireland, M., Gonzales, A., & Booth, R. (2007). *The development and psychometric properties of LIWC 2007* (www.LIWC.net, Austin, TX).

Pennebaker, J., & Francis, M. E. (1996). Cognitive, emotional, and language processes in disclosure. *Cognition and Emotion, 10*(6), 601–626.

Pennebaker, J., & King, L. A. (1999). Linguistic styles: Language use as an individual difference. *Journal of Personality and Social Psychology, 7*(6), 1296–1312.

Ritter, R. S., Preston, J. L., & Hernandez, I. (2014). Happy tweets: Christians are happier, more socially connected, and less analytical than atheists on Twitter. *Social Psychological and Personality Science, 5*(2), 243–249.

Sap, M., Park, G., Eichstaedt, J., Kern, M., Ungar, L., & Schwartz, H. A. (2014). Developing age and gender predictive lexica over social media. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP*, 2014.

Schnabel, L. (2015). Review of why are women more religious than men. *Secularism and Nonreligion, 4*(1), Article 7.

Schwartz, H. A., Eichstaedt, J. C., Kern, M. L., Dziurzynski, L., Lucas, R. E., Agrawal, M., … Ungar, L. H. (2013a). Characterizing geographic variation in wellbeing using Tweets. In *Proceedings of the 7th International AAAI Conference on Weblogs and Social Media (ICWSM)*. Boston, MA.

Schwartz, H. A., Eichstaedt, J. C., Kern, M. L., Dziurzynski, L., Ramones, S. M., Agrawal, M., … Ungar, L. H. (2013b). Personality, gender, and age in the language of social media: The open-vocabulary approach. *PLoS ONE, 8*(9), e73791.

Seligman, M. E. P. (2011). *Flourish: A visionary new understanding of happiness and wellbeing*. New York, NY: Free Press.

Smith, T. B., McCullough, M. E., & Poll, J. (2003). Religiousness and depression: Evidence for a main effect and the moderating influence of stressful life events. *Psychological Bulletin, 129*(4), 614–636.

Stark, R. (2002). Physiology and faith: Addressing the "universal" gender difference in religious commitment. *Journal for the Scientific Study of Religion, 41*(3), 495–507.

Sulaiman, R., & Aquil, A. (2015). Growth patterns of social media in Arab gulf states: An analytic study. *Social Networking, 4,* 23–32.

Sumner, C., Byers, A., Boochever, R., & Park, G. (2012). Predicting dark triad personality traits from Twitter usage and a linguistic analysis of Tweets. In *Proceedings of the 11th International Conference on Machine Learning and Applications, ICMLA*.

Sylwester, K., & Purver, M. (2015). Twitter language use reflects psychological differences between democrats and republicans. *PLoS ONE, 10*(9), e0137422.

Sumner, C., Byers, A., & Shearing, M. (2011). Determining personality traits and privacy concerns from Facebook activity. *Black Hat Briefings* (pp. 197–221).

Tausczik, Y., & Pennebaker, J. (2010). The psychological meaning of words: LIWC and computerized text analysis methods. *Journal of Language and Social Psychology, 29*(1), 24–54.

Thomas, J. (2014). *Psychological wellbeing in the Gulf states: The new Arabia Felix*. London, UK: Palgrave Macmillan.

Thomas, J., Al-Marzooqi, A., & Raynor, M. (2012). Marital status and gender as predictors of undergraduate academic performance: A United Arab Emirates context. *Learning and Teaching in the Gulf, 9*(2), 1–9.

Thomas, J., Al-Shehhi, A., & Grey, I. (2018, Accepted). The sacred and the profane: Social media and temporal patterns of religiosity in the United Arab Emirates. *Journal of Contemporary Religion*.

Wang, W., Chen, L., Thirunarayan, K., & Sheth, A. P. (2014, February 15–19). *Cursing in English on Twitter*. Paper presented at the 17th ACM conference on Computer supported cooperative work & social computing, Baltimore, Maryland, USA.

Watson, D., Clark, L. A., & Tellegen, A. (1988). Development and validation of brief measures of positive and negative affect: The PANAS scales. *Journal of Personality and Social Psychology, 54,* 1063–1070.

Würtz, E. (2005). A cross-cultural analysis of websites from high-context cultures and low-context cultures. *Journal of Computer-Mediated Communication, 11*(1), Article 13.

Yang, C., & Srinivasan, P. (2016). Life satisfaction and the pursuit of happiness on Twitter. *PLoS ONE, 11*(3), e0150881.

Yarkoni, T. (2010). Personality in 100,000 Words: A large-scale analysis of personality and word use among bloggers. *Journal of Research in Personality, 44*(3), 363–373.