
ZOLTÁN KMETTY, JÚLIA KOLTAI, ESZTER BOKÁNYI AND
KÁROLY BOZSONYI *

Seasonality Pattern of Suicides in the US – a
Comparative Analysis of a Twitter Based Bad-mood
Index and Committed Suicides

Intersections. EEJSP
3(1): 56-75.
DOI: 10.17356/ieejsp.v3i1.302
<http://intersections.tk.mta.hu>

* [zkmetty@tatk.elte.hu] (ELTE TÁTK, MTA-ELTE-Peripato
Research Group); [koltai.juli@gmail.com] (MTA TK, ELTE TÁTK);
[e.bokanyi@gmail.com] (ELTE TTK); [bozsonyi.karoly.bk@gmail.com]
(KRE BTK)

Abstract

Is it possible that a general negative social climate exists, which, on the one hand, manifests itself in the number of suicides committed, and on the other hand, appears in the content of tweets posted on social media? In this paper, we are attempting to identify the seasonal and weekly patterns of suicide-related tweets and compare them with similar data that shows the ratios of committed suicides in the US. For this work we used the data stream freely provided by the online social networking site Twitter to collect geo-located tweets in the US.

In order to calculate the negative mood Twitter-index three terms (words) were used related to suicide (suicide) and bad mood (depression, depressed). The raw daily occurrences of these three words were summarised and then divided by the number of all geo-located daily tweets in the US collected in the framework of the project.

The weekly fluctuation of the temporal distribution of suicides and the ratio of bad mood messages on Twitter fit together well. Tweets show a deterioration of mood on Sundays more intensely; this tendency is more moderate in suicide data. Monthly data, however, are much more challenging to interpret, since the fluctuations show a completely opposite tendency than we expected. We present two possible explanation for this unexpected result.

Keywords: Suicide; Seasonality; Bad-mood; Big Data; Twitter.

1. Introduction

The use of Big Data creates numerous new opportunities to analyse and to understand the mechanisms of the social world (Lazer et al., 2009). However, there are also many questions about the utilisation of this kind of data. In most cases (especially when using online data), the lack of external validity is one of the main concerns (Lewis, 2015). In a field as young as this one, it is useful to compare the results with other well-known social trends: if the results based on this kind of data are comparable with the earlier results based on traditional techniques that have more external validity, we can get closer to understanding the validity of the data source better. In this paper, we are attempting to test this kind of validity using a case of suicide-analysis in the USA, in particular by comparing data about the population with data obtained from the online social networking site Twitter (which we can consider as Big Data). It is important to note that this study is not the first in connection with this topic.

A paper from 2014 (Jashinsky et al., 2014) analysed the spatial correlation of suicide and tweets about suicide. They found that the emotional charge of Twitter messages and the suicide rates in the population show significantly similar spatial patterns in the territory of the United States. We are looking for the same correlation, but in the variation of time. Necessarily we do not suppose, that the people who commit suicide, are the same people who show negative emotions on Twitter, we only hypothesise that there is a general social and emotional climate that changes in space and time. This climate affects both Twitter messages and the number of suicides. Therefore, we do not think that the correlation between the two phenomena implies a cause and effect relationship, but we applied a causality-scheme, in which there is a common cause of the emotions that appear in the Twitter data and also of the number of suicides.

It is a well-known phenomenon in suicide research that the incidents of suicide follow a clear seasonal pattern. There are minor differences between countries, but the main trend is that the ratio of suicides is high in the spring and summer, and low in the winter (Warren, Smith and Tyler, 1983; Massing and Angermeyer, 1985; Corcoran et al., 20014; Zonda, Bozsonyi and Veres, 2005). There is also a strong weekly pattern: the risk of suicide is high on Mondays and much lower at the weekends (Massing and Angermeyer, 1985). In this paper, we are attempting to identify the seasonal and weekly patterns of suicide-related tweets. We are using geo-located tweets that have been collected by the Department of Physics of Complex System of Eötvös Lorand University, Budapest since 2012. Using this data, we can compute the seasonal and weekly patterns of suicide-related tweets and compare them with similar data that shows the ratios of suicides committed in the US.

2. Twitter in Social Sciences

Several studies have been published in recent years where data from Twitter was used in different social science analyses. Some of these studies used Big Data techniques while others did not. We would like to review a few important studies which have used the Big Data approach below, as we believe they illustrate the applicability of Twitter data for the purposes of social sciences quite well. These studies involve different research areas, but those conducted about health behaviour, public opinion and social network aspects appear to be the most prominent.

Paul and Drezde (2011) analysed more than 1.5 million health-related tweets. The temporal and spatial distribution of illness and symptom-related posts had strong correlations with the population data published by health authorities. Having analysed the posts of 210 000 Twitter users, Abbar, Mejova, and Weber (2015) found that daily calorie intake figures calculated on the basis of meal-related posts have 0.77 correlation with obesity index figures broken down by states in the US. With the help of data from Twitter, it was possible to come up with a better model for obesity and diabetes distributions at the level of counties. Mitchell et al. (2013) also published analyses in connection with lifestyle; they collected 80 million words of Twitter posts from the US. They examined the co-occurrence and correlations of words with the help of large-resolution spatial databases. They discovered interesting relationships between the contents of Twitter posts and the lifestyles of citizens living in different region of the US.

Dodds et al. (2011) developed an alternative survey method drawing exclusively on Twitter data. They argue that their results are in line with those deriving from traditional survey methods. Jahanbakhsh and Moon's (2014) study should also be mentioned in connection with survey methodologies, as having analysed 32 million Twitter messages before the 2012 presidential election they found that Obama's popularity trend and his expected victory could be accurately traced and predicted based on Twitter data.

There are two recently published studies, which are excellent examples of social network analyses. These analyses are particularly important with regard to the current study because they are based on the same Twitter database that has been used in our study. Szüle et al. (2014) replicated one of the most famous experiments of sociology, Milgram's small world project, on a social graph generated from the network of Twitter users. Having analysed the network of 6 million geo-located Twitter users, they found that in general any user can be reached in 3 to 6 steps starting from any other user. It is a novel finding, however, that the topology of networks within and across cities differs significantly. Kallus et al. (2015) also published a study on social networks, where the authors analysed the networks of 5.8 million geo-referenced Twitter users on a global scale. Their most important finding was that despite the fact that the network is global, actual relationships emulated the regional structure (administrative units) of societies.

As the above-mentioned examples show it clearly, the use of Big Data deriving from social media content is increasingly popular in social science. This phenomenon is reflected by the fact that in 2014 a new journal was released by Sage, titled *Big Data and Society*. The handling and analysis of this kind of data and also the opportunities

offered by it are different from classical survey data, which is what is usually employed by social scientists. National or international representative samples are believed to have high external validity¹. The generalisation of the results from Twitter or other social media analyses is more problematic. Although there are scientific papers that try to investigate the user-population of these online applications (Pennacchiotti and Popescu, 2011; Sloan et al., 2015; Culotta, Ravi and Cutler, 2015), we do not have precise knowledge about the composition of Twitter users and far less about those who make their profile public and are, therefore, researchable². However, the internal validity of Twitter and other social media data is higher, as this data was not created for research; research is only an additional possibility. Intentional distortions in the answers or the effect of the interviewer do not pose problems during the analysis; we get more honest data. Nevertheless, because of the same reason – it is not primarily created for research – the data is not clear, it is almost always noisy, so researchers have to put more effort into the cleaning process. However, the most important aspect for our paper is the role of time in this data. When analysing (inter)national representative surveys, researchers mostly work with cross-sectional data. Even longitudinal studies have limited time-points, so time can only be handled as a discrete variable in most social research. On the contrary, we can collect Twitter data every second, so it is possible to treat time as a real continuous variable and examine trends and returning patterns. As we will discuss later, we would like to detect a climate that has an effect on people's moods, and as we know that there are seasonality and long-term trends in this climate, the time-based opportunities of Twitter data are absolutely essential in our research.

Although we have emphasised the differences between big social data and survey data, it is equally important to note that the methods and concepts of analysis are not radically different compared to the ones used on survey data. Of course, because of the high number of cases and low external validity, significance tests should be used with caution³. Nevertheless, aside from some statistical and technical issues, the concept of the analysis is not radically different. Moreover, it is essential to keep in mind social scientists' conceptual approach when they are analysing massive datasets. Although the statement 'it does not matter why it works, it only matters that it works' is popular among data scientists, the ability to ask good questions and the confirmatory analyses of theories are necessary if we do not only want to understand how social processes work, but we also want to explain them. Thus, this new type of data does not annul previous knowledge but subjects it to new ways of applications and challenges. In the latter part of this paper, we also use classical methodological concepts, for example Lazarsfeld's theory of indicators, correlation analysis, ANOVA, etc.

¹ It is important to keep in mind that refusing to answer and substitutions also make the sample and external validity biased.

² The reliability of the data in the profiles is also questionable in some cases.

³ If we want to know whether educated people earn less or not, this cannot be tested validly on Twitter data using the significance levels of statistical tests because we do not know the population to whom the results of these tests apply.

Because of the reasons above, Twitter data is not really suitable for the exact measurement of social phenomena or making predictions (Gayo-Avello, 2012; Lazer et al., 2014), but it can be considered more of a measurement of latent mood.

3. Seasonal pattern of suicide

The change of seasons induces different organic, behavioural, and psychological changes in people living in the temperate zone. These changes are, on the one hand, adaptive/maladaptive responses of the organism, in particular of the nervous system, to the environmental changes brought about by the change of seasons. On the other hand, as a result of the seasonal embeddedness of social life (what Emil Durkheim called changes in the dynamic density of society), seasonal changes also have an effect on human behaviour, besides psycho-somatic factors, through different social factors.

It is an old observation that self-destructive deeds reach their peak in the springtime and early summer months, while the lowest number of suicides are committed in the late autumn and winter months. This phenomenon can be witnessed in both hemispheres of the globe (Warren, Smith and Tyler, 1983; Massing and Angermeyer, 1985; Corcoran et al., 20014; Zonda, Bozsonyi and Veres, 2005).

There are essentially two paradigms that aim to explain the distinct yearly seasonal fluctuation in the incidence of suicides. Representatives of the neuro-psychobiological approach believe that the biological effects of certain environmental parameters (sunshine, temperature, air pressure, etc.) resulting from the change of seasons might be held responsible for the seasonal changes in suicidal behaviour. The socio-demographic approach, on the other hand, disputes the exclusive validity of the biological explanation on the grounds that seasonal fluctuation is strongly influenced by several psycho-social and social variables, for example, age, gender, settlement type, social status, and days of the week. The significant effect of the different days of the week is particularly interesting (the number of suicides is high on Monday and low over the weekend), as there are obviously no systematic changes in atmospheric-environmental parameters linked to the days of the week. Social life, and thus the dynamics of social relations, are, however, significantly influenced by the days of the week. Therefore, it can be hypothesised that the systematic changes observable in the number of suicides during the week are caused by weekly fluctuations in the dynamics of social relations. These social relations and also reflections on social reality should be tangible through various social network channels. In the next section, we will focus on this particular aspect.

4. Suicide and bad-mood in tweets - conceptual background

We have so far reviewed Big Data studies that have been published in the field of social sciences in recent years. We drew attention to those aspects of validity that need to be addressed when analysing Twitter data and suggested social phenomena that might be in the background of the seasonal fluctuation of suicides. In the next section, we are going to shed light on how Twitter-based sentiment analysis could be linked with the temporal dynamics of suicides.

In our analysis, we looked for Twitter posts with geographical reference to the United States, which contain words referring to negative emotional states and/or suicide/depression.

The population prevalence of suicide is fortunately approximately only a dozen people for 100 000 inhabitants; therefore, it is quite unlikely that we would come across the Twitter post of a person planning actual suicide. Nevertheless, we have reasons to believe that the general emotional climate, which in certain cases might induce suicidal behaviour among those susceptible, probably also affects those not planning to commit suicide. Therefore, its indirect effect might be detectable in the emotional load of Twitter posts (in connection with a similar mechanism see Eichstaedt et al., 2015).

This negative affective conditioning must be particularly strong if it is to appear on social media surfaces since it is a well-known fact that social media users are likely to present a picture of themselves and their environment which is more positive than reality (Stenros, Paavilainen and Kinnunen, 2011). Social media, therefore, act as spontaneous filters, which only allow the display of the effects of the negative emotional-mood climate if this exceeds a certain threshold. Therefore, it can be expected that the lowest points in the fluctuation of collective emotions will also be present in Twitter users' posts.

Beyond the proposed theoretical and principle-based arguments, there are two more specific indications that suggest Twitter messages expressing negative emotional states probably display a pattern similar to that of suicides committed.

In their study, Jashinsky et al. (2014) examined the spatial distribution of Twitter messages with negative emotional loading within the US and compared it to the spatial distribution of suicides, and they found significant correlation between the two. Based on the similarity of the spatial patterning of the two distributions, we are also expecting to find similarities in temporal fluctuation. We are hoping to reveal similar temporal fluctuation in our present study. Cody et al. (2016) in their article titled *Public opinion polling with Twitter* examined the relative frequency of several thousands of words appearing on Twitter over a period of seven years. The monthly relative frequency of the word 'happy' over the seven years examined displays the reverse tendency of the aggregated monthly occurrence of suicides committed. Based on this result, it can be assumed that words related to negative emotional states might display a seasonal pattern that is in line with the occurrence of suicides.

5. Data and methods

We used the data stream freely provided by the online social networking site Twitter through their Application Program Interface (API) that allows the downloading of approximately 1 per cent of all publicly sent messages (tweets).⁴ In this study, we focused on that part of the data stream where the actual location of the user is indicated clearly. These so-called geo-located tweets originate from users who chose


⁴ For the 1% rate of Twitter messages, see one Twitter employee's comment on the official twitter community forum: <https://twittercommunity.com/t/is-there-a-limit-to-the-amount-of-data-the-streaming-api-will-send-out/8482> Accessed: 25-03-2017.

to allow their mobile phones to post their precise GPS coordinates along with a tweet. The total geo-located content was found to comprise only a small percentage of all tweets; therefore, by applying data collection focussing only on these, a large fraction of all geo-tagged tweets can be acquired (Morstatter, Pfeffer and Liu, 2013).

More than 626 million geo-located tweets were collected between 2012 February and 2016 August from the United States of America. Tweets were marked as coming from the USA if their longitude fell between -130 and -70 degrees and their latitude between 24 and 52 degrees. The messages and their metadata were organised into a large relational database that enabled fast and efficient querying (Dobos, 2013) at the Department of Physics of Complex Systems of Eotvos Lorand University, Budapest.

Although this is a large amount of data, the daily frequency of tweets was far from uniform. The number of sent tweets in the US increased over the given period of time, but the number of collected tweets decreased (see Figure 1). One reason for this might be a change in the algorithms of the Twitter API, which was used for the data collection. The exact algorithm is hidden from the users, but it is known that Twitter sometimes makes changes in the tweet sampling processes (Felt, 2016). The applied data collection design has also changed over the period at the university. In order to enable them to answer special research questions, the researchers altered the technical details of the downloading process. Besides this, data collection on ELTE servers also stopped many times because of different reasons (e.g. blackouts in the electricity system).

Figure 1. Number of geo-located tweets in the US on a daily basis



To handle the latter problem, all the days where the number of collected geo-located tweets was under 35 000⁵ were omitted from our analysis. Consequently, although the entire time period from February 1st, 2012 to August 31st, 2016 contained 1705 days, 514 days were dropped due to this reason.

In order to calculate the negative mood Twitter-index, suitable and reliable indicators had to be found. Initially five terms (words) were selected to create the index, these were related to suicide (suicide), bad mood (depression, depressed), and pills used to alleviate depression (Prozac, Zoloft). In this selection, we relied on Jashinsky et al.'s (2014) paper, but instead of searching for long terms, we focussed on unique words.

Table 1 List of search words, and stop words applied

Search term	Stop words	Initial number of tweets	Used number of tweets
suicide	Attack, bomb, bomber, squad, Williams	81638	66538
depression	Economic, concert	59770	52689
depressed		142781	137442
Prozac	LilMiss, Bananas, Clause	1747	1614
Zoloft		672	652

Naturally, not all the filtered tweets related to bad/negative mood, so our data was quite noisy; therefore, this noise level had to be decreased. For this, we attempted to identify stop words that signal different meanings in tweets. We created word-clouds from all tweet contents based on the frequency of the words in all of the selected tweets. Then, in the case of every selected word, we tried to identify unusual words in the word-clouds. Sometimes these stop-words were obvious (like 'bomb' in the case of 'suicide'), but there were also other surprising findings⁶.

Another problem that had to be handled before starting the analysis was that not all the users were persons; institutions can also use Twitter, and some of these institutions tweet very often. Therefore, all the users who had more than 10 tweets per used search words were omitted from our further calculations. After these omissions our data was still noisy, but the quality was much better⁷.

In the next step the data was aggregated at the daily level, so the occurrence of any given word in the tweets was counted per day. As we pointed it out earlier, the data generating and collecting algorithms were changed many times over the data

⁵ The threshold level applied was 10% of the number of average daily tweets.

⁶ 'Clause' (as in Santa Clause) could be cited as one of the interesting findings since it is not straightforward how this word relates to Prozac. The solution is probably quite obvious for those who are familiar with the famous 'Friends' TV sitcom, as there is a scene in it when Phoebe (one of the main characters) ends a discussion with the following sentence 'You are like Santa Claus on Prozac, at Disneyland, getting laid.'

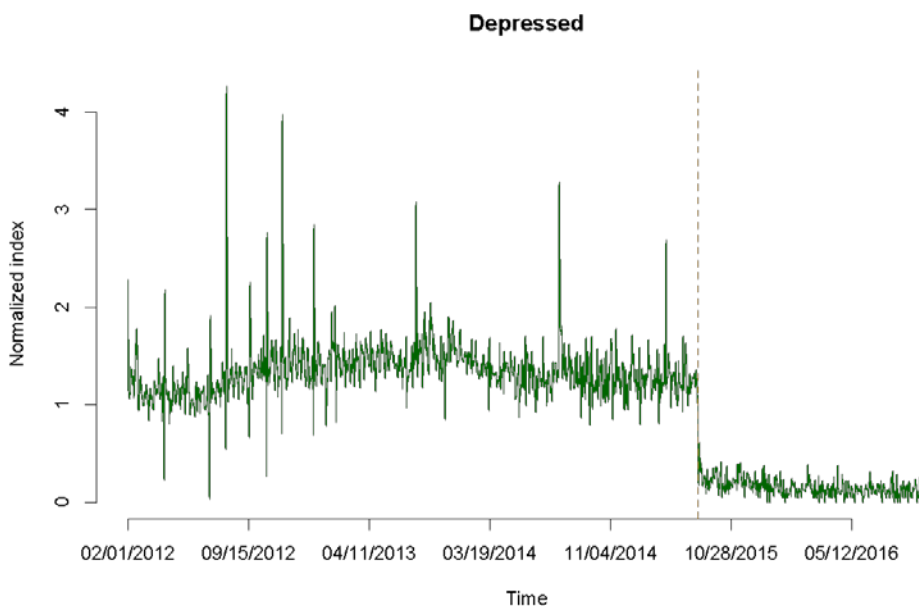
⁷ Obviously not all the tweets were related to suicidal thoughts, but there were many which were exactly tweets about suicide. Here are some examples:

'I just found out one of my old friends tried to commit suicide I feel sick', or 'Strong sense of suicide, I really wanna quit.', or 'Planning my suicide, writing my letter. #whocares #nobody #endingthispointlesslife', 'I'm depressed and I've learned to accept it. Goodnight world.', 'I've never been this depressed everything's falling apart.'

collection period. To handle this issue, the number of daily tweets for each word was divided by the total number of geo-located daily tweets in the US collected in the project. After this normalisation, the average daily tweet number was fitted to one. So for every day, we had the relative importance of the words, compared to the whole US geo-located dataset. This latter standardisation method might help in the comparison of the different search terms. However, the analysis of normalised indices revealed some further problems.

As could be seen in the time sequence of the normalised ‘depressed’ search term, there were some very high peaks in the given period. These peaks usually signal a special event or a unique happening in society. This could be a film premiere (e.g.: *Suicide Squad*), a death (suicide) of a famous person⁸ (e.g.: *Robin Williams*), or a suicide in a TV-show. As mentioned earlier, we had identified some stop words in the initial phase of the project to handle these special events during the period. Despite our efforts, this filtering was not perfect; some peak days still remained in the time series. To increase the reliability of our indices, all these outlier days were filtered out from our analysis⁹.

Figure 2. Normalised index of the occurrence of the word ‘depressed’ on a daily basis



⁸ There is a tendency, that after the suicide of famous people, suicide rate increases due to imitation (Philips, 1974). But as we would like to track the general trends and patterns and not the unique ones, we decided to omit these days from the rest of the analysis.

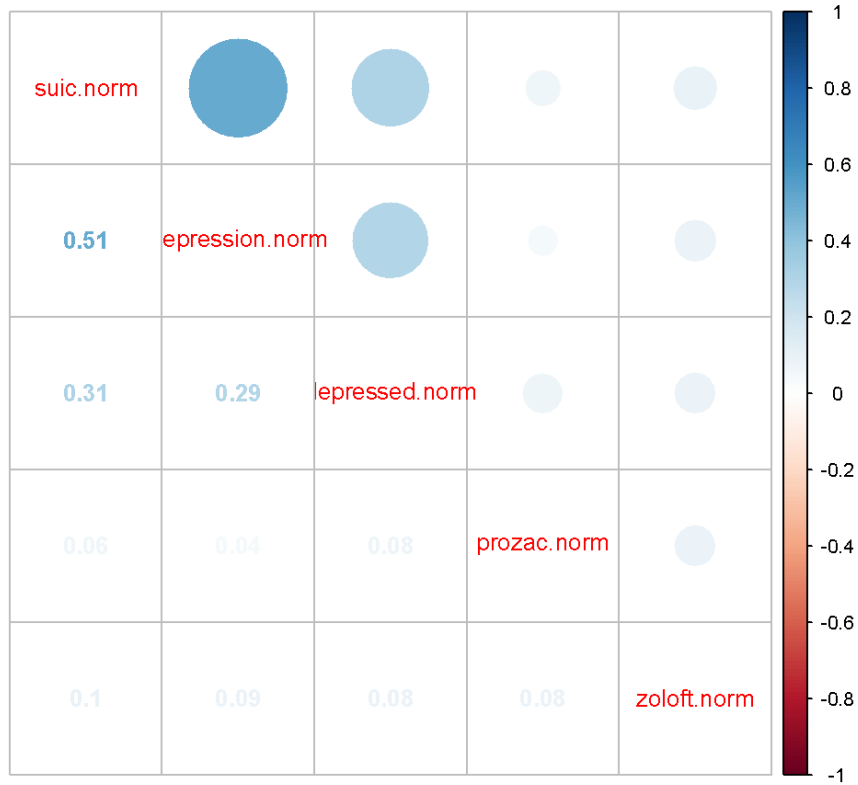
⁹ We filtered out all those days where any of the normalised indices were higher than 3.

There was another anomaly in the time sequence. The mean level of daily normalised tweets suddenly dropped in the case of all indices. The volume of decrease varied between the indices (e.g.: it was higher in the case of 'depressed' than in the case of 'suicide' - see Figure 6 in the appendix). This change could be hard to explain, as these indices were all normalised to the number of daily tweets. The turning point is April 2015, so we needed to search for explanations within this time period. We hypothesised that this could be related to a change in the API, or also to a change in the data collection process. In the end, we decided to disregard all the days after this event, as the process of generating the available tweets probably differs in the two time periods. Finally, 838 days remained in our analysis. During the above data validation processes, nearly 20 percent of the tweets containing the word 'suicide' were filtered out, but this rate was only around 4 percent in the case of the word 'depressed'.

If we assume that all the five search words are related to the same phenomenon - bad mood - then the daily frequency of these indices has to be strongly correlated with each other.

The correlation matrix of the five normalised time series (Figure 3) clearly shows that the daily distributions of the words 'Prozac' and 'Zolofit' are independent of the others. This suggests that these terms were too rare to be included in our calculations. Therefore, we decided to leave out these two indices from our further analysis. The correlation between the other three variables was significant: strong between 'suicide' and 'depression', and average between 'depressed' and the other two words. In order to create our negative mood index, the raw daily occurrences of the three words were summed and then divided by the number of all geo-located daily tweets in the US collected in the framework of the project. The bad-mood index was also balanced to mean equal to one, in order to assist further interpretation.

Figure 3 - (Graphical) Correlation matrix of the normalised occurrence of the five search words on a daily basis



In order to calculate statistics for the suicides committed in the US, we used the CDC WONDER Online database¹⁰. The database contains information about multiple causes of death between 1999 and 2014. We filtered the data to include death caused by suicide only. In some cases the day of the suicide was missing, so we omitted those cases from our further analysis. Over these 15 years, 560 000 suicides were committed in the US, and the yearly suicide rate increased by nearly 30 percent between 1999 and 2014.

For reasons of anonymity, the exact number of suicides committed daily cannot be extracted from the database; therefore, we needed to use data aggregated by days of the week. Thus, we have no precise knowledge about the exact number of suicides on 7 February 2002 for example, but we know the total number of suicides in February 2002 for each day of the week (e.g.: the total number of suicides on Wednesdays in February 2002). In order to suit our analytic aims, this data form needed to be

¹⁰ Centers for Disease Control and Prevention, National Center for Health Statistics. Underlying Cause of Death 1999-2014 on CDC WONDER Online Database, released 2015. Data derives from the Multiple Cause of Death Files, 1999-2014, as compiled from data provided by the 57 vital statistics jurisdictions through the Vital Statistics Cooperative Program. Available at: <http://wonder.cdc.gov/ucd-icd10.html> Accessed: 18-08-2016.

normalised. We created two types of databases, one for the days of the week, and one for the months. In the case of the days of the week, the total number of suicides was divided by the number of that specific day of the week in that particular month. As there was a clear increasing tendency in the number of committed suicides, the data also had to be normalised with the yearly trend. In the case of months, the data was aggregated to months on a yearly basis (using the sum function), and this number was divided by the number of days in that month. This data also needed to be normalised by the yearly trend of suicide.¹¹ As the usual Twitter users are probably younger than the average US population we also calculated the above mentioned suicide statistics for the 15-44 age cohort. As the days of the week and monthly distributions of suicides did not differ in the case of the 15-44-year-olds and the entire population, the 15-44 subpopulation data was not used in further parts of the paper.

6. Results

6.1 Days of the week

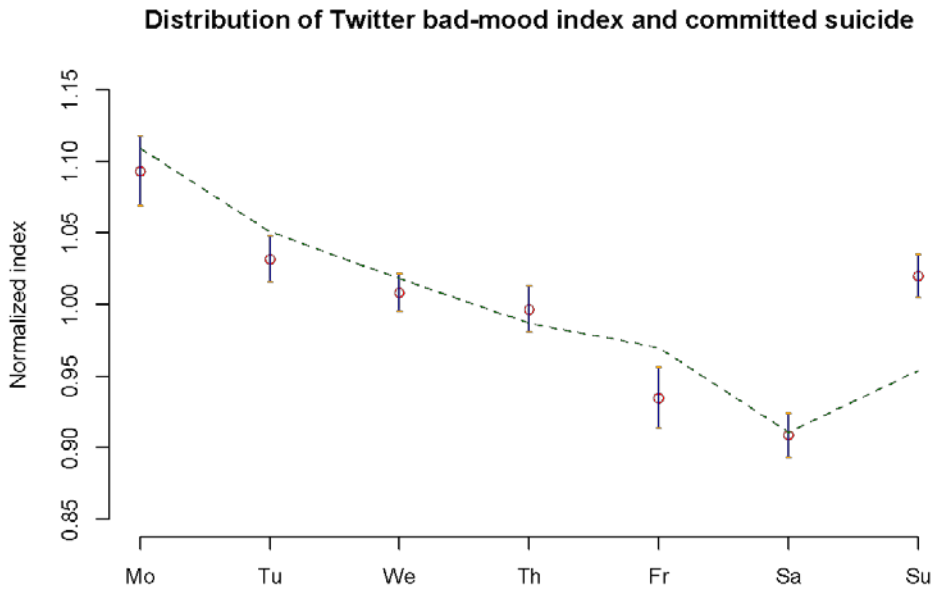
When analysing suicides committed between 1999 and 2014 in the US, a weekly fluctuation is easily identifiable. On average, most suicides were committed on Mondays, which is followed by a decreasing tendency reaching its lowest on Saturday. Thus, the number of suicides decreases as Monday is getting more distant and the weekend is approaching. On Sundays a slight upward correction can be witnessed, that is the suicide rate on Sundays is slightly higher than on Saturdays, but is it still lower than the figures measured on Fridays. The difference between the days is significant ($p=0.00$ on ANOVA tests)¹², and the difference between the highest and lowest figures is 21 percent.

The negative emotional climate identified on the basis of Twitter data displays a very similar distribution over the course of the week. From the high value characteristic of Mondays, the number of tweets suggesting bad mood decreases significantly by Tuesday and stagnates at this lower level until Thursday. On Friday, another change can be witnessed in the emotional climate, the number of bad mood tweets decreases even more, reaching its lowest on Saturday (although the difference between Saturdays and Sundays is not statistically significant). On Sunday, a rebound can be observed in the Twitter negative mood index; the number of negative tweets jumps to the level characteristic of the middle of the week. On the next figure, a dotted line indicates the temporal fluctuation of actual suicide data, while error bars show the daily averages of the bad-mood index and their confidence intervals.

¹¹ This normalisation was done, because when we analysed the Twitter data, we used the same methodology. For that very reason, the normalisation of the population data was necessary for the comparison of the two datasets. (We also checked the weekly and monthly population trends without the normalisation procedure, and we found that the results were similar to the ones that we present in this paper.)

¹² From a theoretical perspective, in the case of twitter data it is appropriate to use significance tests. In the case of suicide data the situation is more complicated, as we have full data, not just a sample. But from Bayesian perspective it is an acceptable viewpoint to calculate significance values. On the other hand, because of the large sample size, we used higher significance-levels in the analysis: instead of the general 0.05, we applied a level of 0.001.

Figure 4 - The weekly fluctuation of suicides committed in the US (green, dotted line), and the variation of the bad-mood index (error bar)




The figure displays the similarity of the two temporal distributions well. The most remarkable difference can be observed in the values for Sundays. The mood of average American Twitter users is already starting to deteriorate by Sunday based on the words analysed, but this can hardly be witnessed in actual suicide rates.

Although the statistical relevance of correlation is quite limited when calculated on such a low number of cases ($N=7$); for the sake of illustration, we would like to note that the correlation between the figures was 0.87, which means that there is a strong and positive relationship between the social network and the population data.

6.2 Months

The monthly fluctuation of committed suicides follows the trend that can be expected based on the literature, which we have already described above¹³. The index starts from the base value of under one in January, and it decreases ever further in February (it is hypothesised that Valentine's day might contribute to this effect¹⁴). After this, the number of suicides starts to increase, and it stagnates from April until July at a relatively high value. Afterward that, the number of suicides starts to decrease only slowly at first, then quite rapidly, reaching its lowest value in December. Seasonality over the months is less marked than within the week, the difference between the maximum and minimum values is 'only' 14 percent.

Figure 5 - The monthly fluctuation of committed suicides in the US (green, dotted line), and the variation of the bad-mood index (error bar)



The fluctuations in bad mood climate reconstructed on the basis of Twitter data seem to be moving in opposite directions from suicides. Winter months (from November till January) can be characterised by relatively more frequent negatively loaded tweets (November is especially outstanding). In the months of transition (February, September, October) negative tweets are relatively fewer (but their number

¹³ Some researchers investigating the relationship between suicides and antidepressants argue that taking medication could weaken or might even eliminate seasonality from suicide time series (Sebestyen et al., 2010). American data provide no support for this hypothesis, what is more, between 2012 and 2014 findings suggest even stronger seasonality effects.

¹⁴ In connection with holidays see also: Zonda et al., 2009.

is still above average), while in the other months, and especially in the middle of summer, this figure reaches its minimum. The correlation between the fluctuations of the two indices is -0.78.

7. Discussion

In our study, we attempted to find answer to the following research question: Is it possible that a general negative social climate exists, which, on the one hand, manifests itself in the number of suicides committed, and on the other hand, appears in the content of tweets posted on social media? Researchers analysing the spatial aspects of this question (Jashinsky et al., 2014) came to the conclusion that it is possible to hypothesise and identify such a climate. In our current study, we attempted to grasp the dynamic, temporal aspect of this question, but our results fail to provide a straightforward answer. The weekly fluctuation of the temporal distribution of suicides and the ratio of bad mood messages on Twitter fit together well. Tweets show a deterioration of mood on Sundays more intensely; this tendency is more moderate in suicide data. Since it can be assumed that accumulated feelings of depression are in the background of many cases of suicide, we believe that this temporal delay does not contradict our hypothesis.

Monthly data, however, are much more challenging to interpret, since the fluctuations show a completely opposite tendency than what have expected. We can put forward two hypotheses based on our results.

1. As has already been pointed out in the section discussing the seasonality of suicides, researchers basically hypothesise two drivers in the background of the seasonal fluctuations. One is a (neuro-) biological explanation, and the other is social. These theories are not mutually exclusive; the two explanations can co-exist. In the case of weekly fluctuations, it would be very difficult to posit any biological factor; therefore, only social influences can have an explanatory role there. In this regard, it is quite reassuring that the weekly fluctuation of negative social climate reconstructed on the basis of Twitter and the weekly fluctuation of committed suicides are practically identical. In the case of monthly seasonality, however, it is plausible to posit biological drivers as well (even if their precise mechanisms are not known at the moment, and only hypotheses exist about them). These biological drivers might even override social ones. If we believe that it is only a relatively small fraction of society that is at risk of committing suicide, then we might also think that this biological effect which exerts its influence in spring-summer does not really concern great masses of people ('only' those, for example, who suffer from severe depression). Therefore, no sign of this can appear in the negative climate identified on the basis of Twitter. This line of argumentation also suggests that social drives would influence the natural course of this process in the opposite direction; therefore, biological drives must be very strong to be able to override this influence.

2. Our other hypothesis is related to the general emotional load of tweets. In an earlier part of our paper, we already cited Cody et al.'s (2016) study, where the temporal fluctuation of the word 'happy' could easily be identified based on the

collected data. Their results showed that Twitter users use the word ‘happy’ more frequently in the winter than in the summer. Based on this article, we expected to find an opposite tendency in connection with the words reflecting bad mood. As our results show, this was not the case; winter was the peak for the words we analysed as well while summer brought lower results. What can be concluded from this? We may hypothesise that the ratio of emotional tweets is not evenly distributed throughout the year. People might be more likely to tweet a higher number of emotion-related messages in general in winter than in summer, regardless whether the emotions are negative or positive. In order to check this assumption, we also analysed the frequency of the word ‘happy’ in the tweets geo-located in the US. The temporal fluctuation of the word ‘happy’ and the negative mood index (we used normalised indices in both cases) had positive correlations above 0.3. This might suggest that in our subsequent analyses we should try to operationalise a further index that would show the parity of positive-negative mood.

Further questions and limitations

It is important to mention that as the explanation of trends in Twitter data, and the correlation with the number of suicides, alternative explanations are possible. However, it is also possible that the cause that affects the mood cached in the tweets and also the number of suicides, is the same cause that affects the frequency of other dimensions in the tweets as well. As we do not have a chance to conduct experiments in this research field, we can never catch real causality, just plain correlations between variables.

Besides the practical interpretation of our findings, it would be worthwhile to consider the lessons that can be drawn from our work in a wider context. Although we touched upon problems of validity in the theoretical section, we had difficulty implementing it in the empirical part of our study. The facts that we have no information on, i.e. the exact demographic composition of Twitter users, which users apply geolocation, or whether the API made accessible by Twitter filters the tweets somehow, all pose challenges to validity. Studies like our current one would like to contribute to this debate on validity by attempting to validate the processes reconstructed based on Twitter along external criteria. Our study also shows that this is not an easy task. Despite the complexity of the task, we believe that there is an increasing need for such studies, adding that in our opinion the most promising direction would be the publication of more, primarily confirmatory, studies. If we can pose sociological questions that can be answered within the Big Data paradigm, then in the long run it will not be a question whether the Big Data approach needs sociologists or not. We hope that our analysis will also be useful in the sense that it demonstrates how sociology might benefit from the Big Data approach.

References

- Abbar, S., Mejova, Y. and I. Weber (2015) *You Tweet What You Eat: Studying Food Consumption Through Twitter*. Paper presented at the 33rd Annual ACM Conference on Human Factors in Computing Systems. April 18-23, Seoul. 3197-3206. DOI: <https://doi.org/10.1145/2702123.2702153>
- Cody, E. M., Reagan, A. J., Dodds, P. S. and C. M. Danforth (2016) Public Opinion Polling with Twitter. Available at: <https://arxiv.org/abs/1608.02024> Accessed: 25-03-2017.
- Culotta, A., Kumar, N. R. and J. Cutler (2015) *Predicting the Demographics of Twitter Users from Website Traffic Data*. Paper presented at the Twenty-Ninth AAAI Conference on Artificial Intelligence. January 25-29, Hyatt Regency, Austin. 72-78. Available at: <https://pdfs.semanticscholar.org/b8f8/980efb2684ba2a0b7ddf8ea88a9fe09f2ade.pdf> Accessed: 25-03-2017.
- Corcoran, P., Reilly, M., Salim, A., Brennan, A., Keeley, H. S. and I. J. Perry (2004) Temporal Variation in Irish Suicide Rates. *Suicide and Life-Threatening Behavior*, 34(4): 429-438. DOI: <https://doi.org/10.1521/suli.34.4.429.53735>
- Dobos, L., Szule, J., Bodnar, T., Hanyecz, T., Sebok, T., Kondor, D., Kallus, Zs., Stéger, J., Csabai, I. and G. Vattay (2013) *A Multi-Terabyte Relational Database for Geo-Tagged Social Network Data*. Paper presented at the 4th IEEE International Conference on Cognitive Infocommunications, CogInfoCom. December 02-05, Budapest. 289-294. DOI: <http://doi.org/10.1109/CogInfoCom.2013.6719259>
- Dodds, P. S., Harris, K. D., Kloumann, I. M., Bliss, C. A. and C. M. Danforth (2011) Temporal Patterns of Happiness and Information in a Global Social Network: Hedonometrics and Twitter. *PloS one*, 6(12): e26752. DOI: <https://doi.org/10.1371/journal.pone.0026752>
- Eichstaedt, J. C. et al. (2015) Psychological Language on Twitter Predicts County-Level Heart Disease Mortality. *Psychological science*, 26(2): 159-169. DOI: <https://doi.org/10.1177/0956797614557867>
- Felt, M. (2016) Social Media and the Social Sciences: How Researchers Employ Big Data Analytics. *Big Data & Society*, 3(1): 2053951716645828. DOI: <https://doi.org/10.1177/2053951716645828>
- Gayo-Avello, D. (2012) *"I Wanted to Predict Elections with Twitter and All I Got Was This Lousy Paper" - A Balanced Survey on Election Prediction Using Twitter Data*. Available at: <http://arxiv.org/abs/1204.6441> Accessed: 25-03-2017.
- Jahanbakhsh, K. and Y. Moon (2014) The Predictive Power of Social Media: On the Predictability of US Presidential Elections Using Twitter. Available at: <https://arxiv.org/abs/1407.0622v1> Accessed: 25-03-2017.

-
- Jashinsky, J., Burton, S. H., Hanson, C. L., West, J., Giraud-Carrier, C., Barnes, M. D. and T. Argyle (2014) Tracking Suicide Risk Factors Through Twitter in the US. *Crisis*, 35(1): 51-59. DOI: <https://doi.org/10.1027/0227-5910/a000234>
- Kallus, Z., Barankai, N., Szüle, J. and G. Vattay (2015) Spatial Fingerprints of Community Structure in Human Interaction Network for an Extensive Set of Large-Scale Regions. *PloS one*, 10(5): e0126713. DOI: <https://doi.org/10.1371/journal.pone.0126713>
- Lazer, D. et al. (2009) Social Science: Computational Social Science. *Science*, 323(5915): 721-723. DOI: <http://doi.org/10.1126/science.1167742>
- Lazer, D., Kennedy, R., King, G. and A. Vespignani (2014) The Parable of Google Flu: Traps in Big Data Analysis. *Science*, 343(6176): 1203-1205. DOI: <https://doi.org/10.1126/science.1248506>
- Lewis, K. (2015) Three Fallacies of Digital Footprints. *Big Data & Society*, 2(2): 1-4. DOI: <https://doi.org/10.1177/2053951715602496>
- Massing, W. and M. C. Angermeyer (1985) The Monthly and Weekly Distribution of Suicide. *Social Science & Medicine*, 21(4): 433-441. DOI: [https://doi.org/10.1016/0277-9536\(85\)90223-0](https://doi.org/10.1016/0277-9536(85)90223-0)
- Mitchell, L., Frank, M. R., Harris, K. D., Dodds, P. S. and C. M. Danforth, C. M. (2013) The Geography of Happiness: Connecting Twitter Sentiment and Expression, Demographics, and Objective Characteristics of Place. *PloS one*, 8(5): e64417. DOI: <https://doi.org/10.1371/journal.pone.0064417>
- Morstatter, F., Pfeffer, J. and H. Liu (2014) *When Is It Biased? Assessing the Representativeness of Twitter's Streaming API*. Paper presented at the 23rd International Conference on World Wide Web. April 07-11, Seoul. 555-556. Available at: <https://arxiv.org/abs/1401.7909v1> Accessed: 25-03-2017.
- Paul, M. J. and M. Dredze (2011) You Are What You Tweet: Analyzing Twitter for Public Health. *ICWSM*, (20): 265-272.
- Pennacchiotti, M. and A. M. Popescu (2011) *Democrats, Republicans and Starbucks Afficionados: User Classification in Twitter*. Paper presented at the 17th ACM SIGKDD international conference on Knowledge discovery and data mining. August 21-24, San Diego, CA. 430-438. DOI: <http://doi.org/10.1145/2020408.2020477>
- Phillips, D. P. (1974) The Influence of Suggestion on Suicide: Substantive and Theoretical Implications of the Werther Effect. *American Sociological Review*, 39(3): 340-354. DOI: <https://doi.org/10.2307/2094294>
- Sebestyen, B. et al. (2010) Gender Differences in Antidepressant Use-Related Seasonality Change in Suicide Mortality in Hungary, 1998-2006. *The World Journal of Biological Psychiatry*, 11(3): 579-585. DOI: <https://doi.org/10.3109/15622970903397722>

- Sloan, L., Morgan, J., Burnap, P. and M. Williams (2015) Who Tweets? Deriving the Demographic Characteristics of Age, Occupation and Social Class from Twitter User Meta-Data. *PloS one*, 10(3): e0115545.
DOI: <http://doi.org/10.1371/journal.pone.0115545>
- Stenros, J., Paavilainen, J. and J. Kinnunen (2011) *Giving Good'face': Playful Performances of Self in Facebook*. Paper presented at the 15th International Academic MindTrek Conference: Envisioning Future Media Environments. September 28-30, Tampere. 153-160. Available at:
<http://dl.acm.org/citation.cfm?id=2181062&dl=ACM&coll=DL&CFID=743158995&CFTOKEN=55567363> Accessed: 25-03-2017.
- Szűle, J., Kondor, D., Dobos, L., Csabai, I. and G. Vattay (2014) Lost in the City: Revisiting Milgram's Experiment in the Age of Social Networks. *PloS one*, 9(11): e111973. DOI: <https://doi.org/10.1371/journal.pone.0111973>
- Warren, C. W., Smith, J. C. and C. W. Tyler (1983) Seasonal Variation in Suicide and Homicide: a Question of Consistency. *Journal of Biosocial Science*, 15(3): 349-356. DOI: <https://doi.org/10.1017/s0021932000014693>
- Zonda, T., Bozsonyi, K. and E. Veres (2005) Seasonal Fluctuation of Suicide in Hungary between 1970–2000. *Archives of Suicide Research*, 9(1): 77-85.
DOI: <https://doi.org/10.1080/13811110590512967>
- Zonda, T., Bozsonyi, K., Veres, E., Lester, D. and M. Frank (2009) The Impact of Holidays on Suicide in Hungary. *OMEGA-Journal of Death and Dying*, 58(2): 153-162. DOI: <https://doi.org/10.2190/om.58.2.e>

Used Database

Centers for Disease Control and Prevention, National Center for Health Statistics. Underlying Cause of Death 1999-2014 on CDC WONDER Online Database, released 2015. Data derives from the Multiple Cause of Death Files, 1999-2014, as compiled from data provided by the 57 vital statistics jurisdictions through the Vital Statistics Cooperative Program. Accessed at <http://wonder.cdc.gov/ucd-icd10.html> on Aug 18, 2016 3:48:13 AM

Appendix

Figure 6. Smoothed ratio of 'suicide' and 'depressed' words on a daily basis

