

# HappyMeter: An Automated System for Real-Time Twitter Sentiment Analysis

Joaquim Perotti Canela  
Department of Computer Science  
Manhattan College  
New York, USA

Tina Tian  
Department of Computer Science  
Manhattan College  
New York, USA

**Abstract**—The paper presents HappyMeter, an automated system for real-time Twitter sentiment analysis. More than 380 million tweets consisting of nearly 30,000 words, almost 6,000 hashtags and over 5,000 user mentioned have been studied. A sentiment model is used to measure the sentiment level of each term in the contiguous United States. The system automatically mines real-time Twitter data and reveals the changing patterns of the public sentiment over an extended period of time. It is possible to compare the public opinions regarding a subject, hashtag or a Twitter user between different states in the U.S. Users may choose to see the overall sentiment level of a term, as well as its sentiment value on a specific day. Real-time results are delivered continuously and visualized through a web-based graphical user interface.

**Keywords**—Twitter; social networks; data mining; sentiment analysis

## I. INTRODUCTION

Twitter has become an increasingly popular microblogging service that allows users to publish messages, a.k.a. tweets [1]. It functions as a platform for people to express themselves, which often carries opinions on different subjects. Twitter usage is growing exponentially. There are 328 million monthly active users on Twitter and over 500 million tweets are created per day [2].

The rapid growth of Twitter and the public access of tweets have made Twitter a popular research subject. For example, researchers have examined the use of Twitter in promoting products and sharing consumer opinions [3]. Enterprises have studied the usefulness of Twitter in organizational communication and information-gathering [4]. Furthermore, tweets have been monitored to detect earthquakes [5].

In this paper, we present HappyMeter, a sentiment analysis tool that measures happiness on Twitter. Sentiment analysis is to computationally categorize opinions expressed in a given text. It is essentially important in social media monitoring as it provides an overview of the public sentiment regarding certain topics.

Unlike other online articles, Twitter messages share several unique features. Firstly, the vernacular on Twitter is informal [6]. There could be misspelled words, slang and acronyms in a tweet due to Twitter's informal language style [7]. Secondly, every tweet has a length constraint of maximum 140 characters [8]. Moreover, Twitter covers an exceedingly broad range of topics [6]. Lastly, due to the wide usage of mobile devices and

the rapid flow of tweets, this user-generated data reflect instant reactions as events evolve. Therefore, we built our system intended for providing real-time insights of the public sentiment and showing changes over time.

Our paper presents an automated sentiment analyzer based on the Twitter traffic. The system streams all the tweets published in the contiguous United States in real time. For each tweet, a sentiment score is computed using a statistical sentiment model and the geographical data associated with the tweet are stored. We developed a web-based graphical user interface to deliver results instantly and continuously.

The rest of the paper is organized as follows. Section II reviews the related work. In Section III, we describe the data set used to build the system and introduce the methods and algorithms adapted in measuring the data. Section IV presents the results of the study and the visualization we have built. Section V concludes the paper and proposes future directions.

## II. RELATED WORK

Applications of sentiment analysis are broad and powerful. As a subfield of Machine Learning and Natural Language Processing, research has been conducted ranging from document level classification [9] to determining the polarity (positive, negative or neutral) of sentences [10] and terms [11]. In recent years, sentiment analysis on Twitter, specifically, has attracted increasing attention from many research communities. For instance, Bollen et al. investigated whether public mood on Twitter is correlated with shifts in the stock market [12]. Vegas et al. modeled the 2016 U.S. presidential campaign in the context of Twitter [13]. Zeitzoff used data on Twitter to measure social movements [14].

To determine the sentiment of a tweet, many past studies have focused on supervised learning where the training data are collected based on emoticons, hashtags or both [15], [16]. Experiments show, however, that they contain biased information in sentiment analysis [7]. Another common practice is to manually annotate the data in order to build a pool of training data. The apparent disadvantage of this method is the intensive labor and time involved in the process.

Our work is inspired by Dodds et al.'s study on temporal patterns of happiness on Twitter, in which they used a corpus mapped with happiness scores to examine the sentiment variations on different expressions over time [17]. The expressions, however, consist of mainly words. Other crucial

elements of a tweet, such as hashtags and user mentions, have not been much studied. Moreover, there was no geographical comparison; for example, exploring the public sentiment of a term between different states in the U.S. A later study by Mitchell et al. considered the geographical factor [18], but the results did not reveal the changing patterns over time, for example, observing the sentiment of a term in a particular state over an extended period of time.

In this paper, we present a system that shows the public sentiment of every term on Twitter, including unigrams, hashtags and user mentions. The system computes the public sentiment of every term in each contiguous U.S. state. The process is repeated daily. Results are visualized to reveal the sentiment alternation over time, as well as the comparisons between different states.

### III. THE SYSTEM

The system performs a sentiment analysis on a Twitter tweet corpus collected since June 2016. In this section, we discuss the data set used in the study and how we define and calculate sentiment.

#### A. Data Set

The Twitter Streaming API [19] allows us to crawl real-time tweets and receive instant updates. Currently, we have gathered over 380 million tweets with geographical annotation enabled from the contiguous United States. This number keeps growing in the rate of 1.4 million tweets per day on average. Table 1 shows the basic statistics of the data set in the study. The highest volume of tweets was received on November 8, 2016, when the United States presidential election took place. The system collected more than 2 million tweets (2,292,345 to be precise) from the contiguous U.S. in a day.

TABLE I. STATISTICS OF THE TWEETS COLLECTED

	Number of Tweets
<b>Total Tweets Collected</b>	388,664,640
<b>Average Daily Tweets Collected</b>	1,423,680
<b>Standard Deviation</b>	199,943

As mentioned in Section II, one of the major contributions of this work is the sentiment mining of some key components on Twitter, such as hashtags and user mentions. Thus far, the extracted tweets consist of a massive corpus of more than 29,000 unique unigrams, almost 6,000 distinctive hashtags and over 5,000 different user mentions. Table 2 provides an up-to-date summary of the individual terms collected in the study. Our system performs a sentiment analysis on each of these terms in the context of Twitter.

TABLE II. STATISTICS OF THE TERMS EXTRACTED

	Number of Terms
<b>Words</b>	29,455
<b>Hashtags (#)</b>	5,991
<b>User Mentions (@)</b>	5,074
<b>Total</b>	40,520

#### B. Defining Sentiment

The sentiment of a term in a Twitter message is determined by an existing sentiment lexicon, the dictionary of Language Assessment by Mechanical Turk [17]. The list contains over

10,000 most popular words with their average sentiment score, ranging from 1 to 9. In general, happy words have a high sentiment value with a score close to 9, while sad words are usually associated with a low score. Table 3 shows a sample of the lexicon.

TABLE III. SAMPLE OF THE SENTIMENT LEXICON

Word	Sentiment Score
laugh	8.22
dancing	7.08
torch	5.11
tension	2.94
suicide	1.30

#### C. Processing

The system collects real-time tweets through the streaming API and saves them on a server for further processing. As seen in Fig. 1, the process includes data manipulations, such as data cleaning and location identification, and sentiment computation. Results are thereafter stored in a database. The rest of this section elaborates the processing procedure.

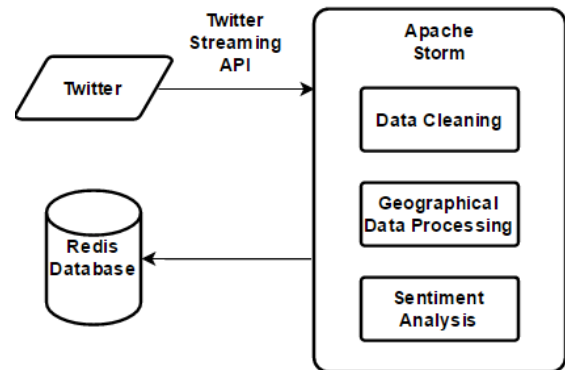


Fig. 1. System architectural overview.

Twitter has the geotagging feature (Tweeting with Location) which allows users to publish a tweet with their location [20]. This feature helps to make tweets more contextual. In the meantime, it provides valuable data for research. One thing to note is that users must give explicit permission for their exact location to be displayed with their tweets, due to Twitter's user privacy policy. Thus, not all the tweets collected come with geographical data. In this study, we keep only the geo-tagged Twitter messages. For each tweet, we store the state where the tweet was issued.

As our work targets tweets in English, tweets written in other languages are discarded. Non-English characters in a tweet are also erased. Due to the informal language model on Twitter (mentioned in Section I), misspelled words can often occur. The system cleans the data by removing words that do not exist in the sentiment lexicon. Hashtags and user mentions, however, are kept in the data set.

To compute the sentiment of a tweet, the system performs a simple average on the sentiment score of each word. Hashtags and user mentions are excluded in this process. Moreover, stop words with a neutral sentiment value falling between 4 and 6 are also excluded in the calculation, following Dodds et al. [17].

Let us take the following tweet for an example, “@missnemmanuel is so gorgeous! #GoTS7e2 #GoT #newcrush”.

Among the terms, “is” has a sentiment score of 5.18, while “so” and “gorgeous” have a sentiment value being 5.08 and 7.42, respectively. Discarding hashtags, user mentions and neutral stop words, only word “gorgeous” is kept in calculating the sentiment of the tweet. Therefore, the average of 7.42, which is 7.42 itself, is assigned to the example tweet.

The system associates this computed sentiment value with every term in the tweet, including hashtags, user mentions and even stop words. Thus, in the previous example, each of the following terms receives a sentiment score of 7.42 along with the tweet: they are @missnemmanuel, is, so, gorgeous, #GoTS7e2, #GoT and #newcrush.

Each of the terms in the example is highly likely to appear in other tweets as well. The system collects the sentiment scores of a term in all occasions in a day and concludes a mean value. For instance, if @missnemmanuel is mentioned 10,000 times in one day, the system would gather the sentiment values from the 10,000 tweets and compute the average. In this work, we examine the daily sentiment of a term in the contiguous United States as a whole, as well as in each state.

One may question the need of computing sentiment of a neutral word. It may not seem necessary from the previously given tweet. But let us consider word “governor” as another example. According to the lexicon, it has a sentiment value of 5.14, which falls in the range of a neutral stop word. However, it would be interesting to see that some states share a higher sentiment value towards “governor” than others do. Moreover, it would be especially interesting to observe the changing pattern over time.

Another concern one may raise is the capacity and scalability of the system. After all, there are millions of potential user mentions and hashtags on Twitter. Plus, new ones are emerging in every second. Keeping a daily record of sentiment for all of them would require tremendous spatial resources and computing power. To tackle this issue, we set a threshold of 3 to be the minimum daily occurrences of a term. Hashtags or usernames mentioned less than three times in a day in the contiguous U.S. are discarded from the database. We believe that this method can help filtering only the active hashtags and popular user mentions.

#### IV. RESULTS

In this section, we demonstrate our system and present results from the analysis. We first studied the amount of daily tweets issued by each contiguous state. To justify the different populations in each state, we calculated the average number of tweets published per day per 10,000 capita in a state. Population estimates are retrieved from the United States Census Bureau [21]. Fig. 2 shows the results after applying Jenks natural breaks optimization [22]. Results ranging from 18 to 80 have been divided into five groups. Among the contiguous U.S. states, Louisiana delivers the most tweets per day per capita, while Wyoming has the least number of daily tweets per capita, as seen in Table 4.

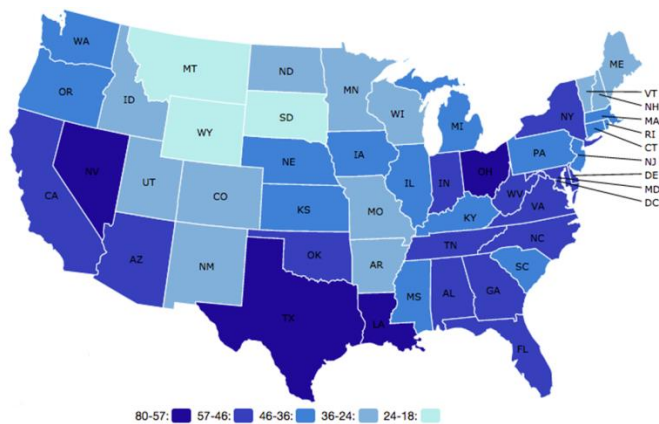


Fig. 2. Average number of daily tweets per 10,000 people in the contiguous United States.

TABLE IV. NUMBER OF DAILY TWEETS PER 10,000 CAPITA

	Daily Tweets per 10,000 Capita	States
<b>Highest</b>	80	Louisiana
<b>Lowest</b>	18	Wyoming
<b>Average</b>	44	N/A
<b>Standard Deviation</b>	13	N/A

Using the methodology introduced in Section III, we investigated the overall sentiment value of each U.S. state based on tweets collected from that region. Results range from 5.92 to 6.03 with a small standard deviation. The average sentiment value for the contiguous United States overall is 5.96. In our study, West Virginia and Wisconsin have the highest average sentiment value, while Alabama, Arkansas, Vermont and Virginia share the lowest sentiment score. Table 5 shows a summary of the statistics.

TABLE V. AVERAGE SENTIMENT SCORES

	Sentiment Score	States
<b>Highest</b>	6.04	Minnesota, Iowa, Nebraska and Utah
<b>Lowest</b>	5.92	Delaware
<b>Average</b>	5.99	N/A
<b>Standard Deviation</b>	0.03	N/A

Similar to Fig. 2, Jenks classification algorithm was utilized to visualize the variation of average sentiment values among U.S. states. Fig. 3 shows the results with five splits.

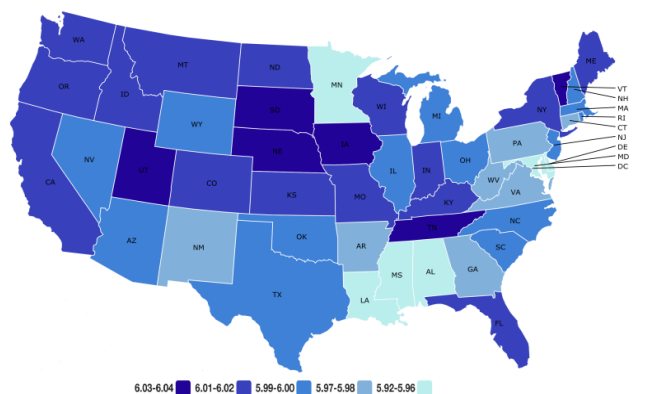


Fig. 3. Average sentiment scores of states in the contiguous United States

This paper also examines the overall sentiment level of each word tweeted in the network. As one can see from the histogram in Fig. 4, most of the words (nearly 75%) have an average sentiment score between 6 and 7, which is considered positive. More than 21% of the words have an overall sentiment value falling between 5 and 6. The highest sentiment score of a word is 8.01 and the lowest sentiment value calculated is 3.73.

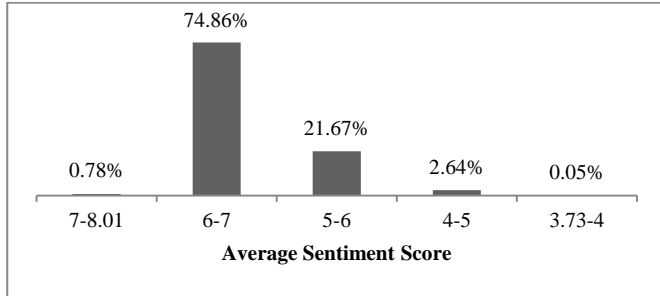


Fig. 4. Histogram of average sentiment scores.

Recall in Section III, a sentiment lexicon built with Language Assessment by Mechanical Turk (LabMT) was used to determine the initial sentiment of a standalone word. Each term was then processed by the HappyMeter system for the overall sentiment in the context of tweets. Table 6 shows the comparisons between the sentiment values before and after the processing of our system. As shown in the table, averagely, the overall sentiment has increased in the Twitter context. Moreover, the standard deviation has significantly dropped, meaning there are fewer extreme ratings. In general, contextual sentiment has become higher and milder. The two sentiment values have a strong Pearson’s correlation of 0.73.

TABLE VI. CHANGE OF SENTIMENT SCORES BEFORE AND AFTER HAPPY METER

	LabMT Sentiment Score	HappyMeter Sentiment Score
Average	5.38	5.97
Standard Deviation	1.08	0.64

TABLE VII. TOP 10 WORDS, HASHTAGS AND USER MENTIONS WITH THE HIGHEST SENTIMENT SCORES

Rank	Word	Score	Hashtag	Score	User Mention	Score
1	birthdayyy	8.01	#beats	7.42	@wwwbigbalthead	6.91
2	brotherly	7.71	#sweepstakes	7.17	@taylorswift13	6.88
3	gday	7.60	#birthday	7.16	@samheughan	6.86
4	belated	7.50	#happybirthday	7.14	@blakeshelton	6.85
5	birthday	7.46	#sweeps	7.13	@jlo	6.85
6	happy	7.41	#shoutout	7.08	@iheartradio	6.81
7	happy	7.38	#flowers	7.07	@britneyspears	6.80
8	splendid	7.35	#beach	7.06	@shawnmendes	6.79
9	congratulations	7.34	#karaoke	7.06	@applebees	6.79
10	unconditional	7.34	#puppylove	7.05	@ethandolan	6.79

TABLE VIII. TOP 10 WORDS, HASHTAGS AND USER MENTIONS WITH THE LOWEST SENTIMENT SCORES

Rank	Word	Score	Hashtag	Score	User Mention	Score
1	headache	3.73	#atxtraffic	3.75	@cnn	5.56
2	killling	3.93	#atltraffic	3.89	@foxnews	5.61
3	kill	3.96	#dfwtraffic	3.89	@thehill	5.62
4	murder	3.98	#traffic	4.08	@msnbc	5.67
5	dead	3.99	#sfltraffic	4.39	@cnnpolitics	5.69
6	accident	4.00	#fail	4.94	@senatemajldr	5.70
7	crying	4.02	#weather	5.16	@nytimes	5.72
8	ouch	4.02	#tampabay	5.36	@nbcnews	5.72
9	jail	4.02	#breaking	5.45	@politico	5.72
10	sick	4.03	#audible	5.69	@senategop	5.72

Table 7 shows the top 10 words, hashtags and user mentions with the highest sentiment level. As we can see from the top words and hashtags, most people feel happy when they tweet about birthdays, flowers, beaches, karaoke and puppies. The best rated Twitter users include mostly movie and television stars, singers and comedians. Table 8, on the other hand, shows a list of the top 10 words, hashtags and user mentions with the lowest sentiment values. As shown in the table, besides the extreme words, such as murder, kill and dead, traffic is the number one problem in people’s common life. User mentions that are associated with low sentiment values are mainly Twitter accounts belonging to the news media.

In this work, we also examined the frequencies of each term on the Twitter network. Table 9 gives a glance of the most often appeared words during the observation, along with their sentiment score. Note that neutral stop words (mentioned in Section III) have been excluded from the list. For example, “just” is the most used word appearing over 12 million times in our data set. However, it is not collected in the list because it can tell little about the public state of mind. We are happy to report that all of the top 10 popular words have a positive sentiment polarity with a sentiment value greater than 6.

TABLE IX. TOP 10 MOST FREQUENTLY APPEARED WORDS (EXCLUDING STOP WORDS)

Rank	Word	Occurrence	Score
1	like	23,880,374	6.38
2	love	16,934,940	7.29
3	see	13,304,971	6.14
4	good	12,727,193	6.59
5	day	12,204,366	6.41
6	lol	11,388,765	6.27
7	great	10,811,384	7.02
8	will	10,507,779	6.02
9	happy	10,177,887	7.41
10	today	9,759,280	6.19

To further investigate the language model on the Twitter network, we built a histogram of usage frequencies of words in the data set, as seen in Fig. 5. Among the 380 million tweets we have received, 35% of the English words appeared between 10,000 and 100,000 times. Interestingly, 26% of the words occurred only less than 100 times. This again proved the casual language style and the rapid change of vocabulary on Twitter.




Fig. 5. Histogram of word frequencies.

We studied the most frequently occurred hashtags, shown in Table X along with their average sentiment value. Job related subjects, such as #job and #hiring, rank at the top of the list leaving the rest far behind. Hashtag #traffic ranks at number 15 with 426,309 mentions and #trump rank at number 18 with 377,559 appearances (not shown in Table 10). Hashtag #education appears later in the list, with 206,122 tags ranking at number 34. Geographical hashtags wise, New York attracts the most attention with 253,696 times mentioning #newyork and 228,345 references of #nyc. Followed after it are #houston with 369,117 occurrences and #chicago with 260,625 tags. The complete list of rankings is available upon request.

TABLE X. TOP 10 MOST FREQUENTLY APPEARED HASHTAGS

Rank	Hashtag	Occurrence	Score
1	#job	17,983,887	6.54
2	#hiring	15,404,646	6.55
3	#careerarc	7,903,993	6.59
4	#jobs	3,537,381	6.51
5	#hospitality	2,568,649	6.57
6	#nursing	1,623,456	6.48
7	#veterans	1,491,045	6.60
8	#retail	1,410,858	6.61
9	#healthcare	1,209,446	6.50
10	#sales	845,894	6.59

Table 11 lists the top 10 most mentioned Twitter users and their overall sentiment score. As we can see, politicians dominate the list. The United States president Donald Trump (@realdonaldtrump) has been quoted more than 2.6 million times during our observation, which is nearly three times more than the second place, Hillary Clinton, his formal presidential campaign competitor. User @potus (President of the United States) ranks at number 4 and Kellyanne Conway (@kellyannepolls) holds the 9th place in the list. The rest of the list consists of mainly television news channels, such as

Fox News (@foxnews), CNN (@cnn), New York Times (@nytimes) and MSNBC (@msnbc). Sean Hannity (@seanhannity), the radio and television host from Fox News, also has been frequently mentioned by the Twitter community, ranking number 10 in the list. The only Twitter account appearing in the top 10 list that is not politics-related is YouTube (@youtube), which holds the 6th place with over 250,000 mentions. The second popular non-political account is @nfl (National Football League), who received less than 100,000 quotes with a rank of 20.

TABLE XI. TOP 10 MOST FREQUENTLY APPEARED USER MENTIONS

Rank	User Mention	Occurrence	Score
1	@realdonaldtrump	2,636,189	5.74
2	@hillaryclinton	707,982	5.75
3	@foxnews	594,875	5.61
4	@potus	533,544	5.83
5	@cnn	511,286	5.56
6	@youtube	256,558	6.17
7	@nytimes	215,899	5.72
8	@msnbc	206,649	5.67
9	@kellyannepolls	190,334	5.80
10	@seanhannity	162,287	5.84

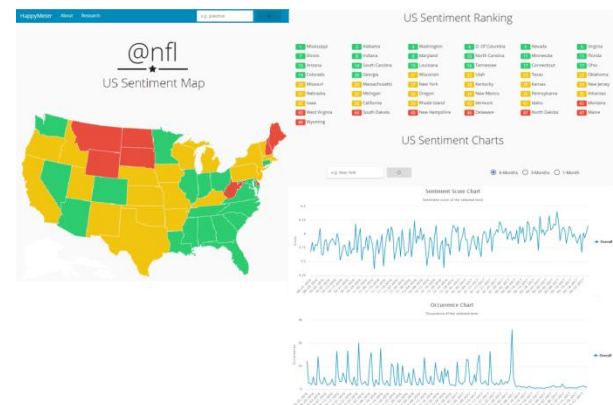


Fig. 6. HappyMeter dashboards of @nfl.

The analysis results of the system are visualized through a web-based graphical user interfaces available at www.happymeter.us. The dashboards display the Twitter sentiment map of a given term, sentiment rankings from the highest to the lowest among states in the contiguous U.S. and charts to reveal the temporal patterns. An example of the dashboards for Twitter user @nfl is shown in Fig. 6.

The sentiment map shows the average sentiment score of a selected term in each contiguous state. To better understand the geography of the public opinions regarding a subject, we applied Jenks natural breaks optimization to cluster the states into three classes. States with higher sentiment scores are classified as the (relatively) positive group. States with lower sentiment values are categorized as the (relatively) negative groups and the remaining states are part of the neutral class. On



the sentiment map, the green color is used to mark the positive group, while states belonging to the neutral and negative clusters are colored with yellow and red, respectively. The map is made interactively to display the sentiment score and polarity of a state when the mouse is hovered over. Fig. 7 shows an example of the sentiment map of Twitter user @realdonaldtrump.

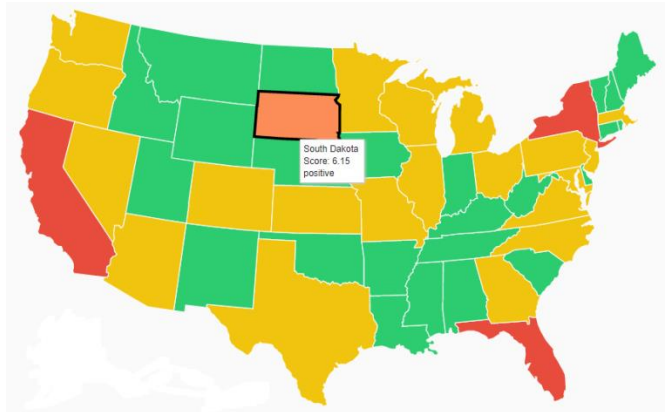


Fig. 7. Twitter sentiment map of @realdonaldtrump.

Besides the sentiment map and rankings, there are two charts on the dashboards. One is the sentiment score chart and the other is the occurrence chart. The sentiment score chart reveals the temporal patterns of a chosen term over an extended period of time. Users are able to choose if they want to see the overall pattern in the contiguous United States as a whole or the sentiment trend in a particular state. Fig. 8 shows an example sentiment score chart of hashtag #job in the state of New York. As one can see, the interactive chart displays the date and sentiment value on that specific day when a data point is selected.

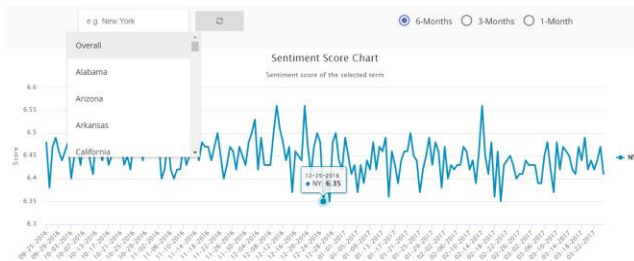


Fig. 8. Twitter sentiment score chart of #job in New York State.

The last component of the dashboards to introduce is the occurrence chart. In addition to sentiment values, the system also keeps track of the daily appearances of a Twitter term. Similar to the sentiment score chart, users are able to project the diagram on the overall contiguous U.S. as well as each state. Fig. 9 shows the Twitter occurrence chart of term “trump” in the whole contiguous U.S. The highest point showing in the figure represents a burst of tweet volume mentioning trump. It occurred on November 8, 2016, the day of the U.S presidential election.

The system uses Apache Storm to gather real-time tweets and manipulate data. Records of sentiment values are stored in a Redis database. The graphical web interface is developed in Python and JavaScript.

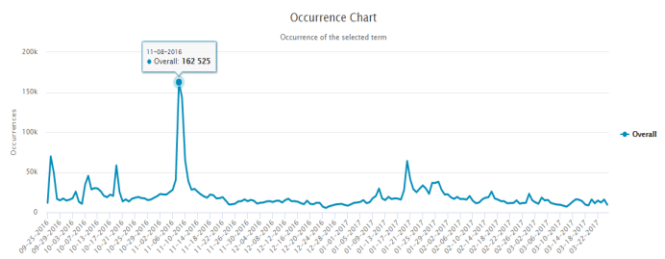


Fig. 9. Twitter occurrence chart of trump.

## V. CONCLUSIONS AND FUTURE WORK

This paper presents HappyMeter, a real-time data processing infrastructure to evaluate public sentiment changes in the context of Twitter. The system examines every term tweeted in the contiguous United States and computes their sentiment scores in the range of 1 through 9. Daily analysis has been conducted throughout the contiguous U.S. as well as in each state. Over 40,000 terms extracted from more than 380 million tweets have been studied. These terms include words, hashtags and user mentions. The system shows the sentiment map and state rankings for each given term. Sentiment charts are automatically generated to reveal the changing pattern of the public sentiment towards a term in the nation or a selected state.

The study also investigates the amount of daily tweets published in a state, as well as its overall sentiment. Interesting findings have been conducted regarding word frequencies, terms with the highest and the lowest sentiment values and the most frequently tweeted words, hashtags and users. The complete analysis results can be made available upon request.

One limitation of the study is that the sentiment lexicon used in the experiment does not cover all the terms. Due to the informal language model on Twitter, new slangs, abbreviations and acronyms are created each day, many of which are Twitter-specific. For example, “twitterati” is a popular term in the Twitter community, which stands for popular users on Twitter. Future work includes designing a mechanism to regularly update the lexicon in order to expand the vocabulary of the dictionary.

Another limitation of the presented work is that context was not taken into account while calculating the sentiment value of a tweet. Our system determines the sentiment by averaging the sentiment score of each unigram. This method performs well in most cases, especially when the data set is at large. But there are times that an average is not able to reflect the true sentiment of a tweet. This is particularly the situation when a sentence is stated as double negative or laid out ironically. In the future, we plan to investigate sentiment scores of n-grams, specifically phrases, in order to achieve results with higher accuracy.

## REFERENCES

- [1] C. C. Aggarwal, Social Network Data Analytics. Springer, 2011.
- [2] Twitter, <https://about.twitter.com/company>, retrieved on 07/28/2017.
- [3] B. J. Jansen, M. Zhang, K. Sobel, and A. Chowdury, “Twitter power: tweets as electronic word of mouth,” Journal of the American Society for Information Science and Technology, vol. 60, pp. 1-20, 2009.

- [4] A. Archambault and J. Grudin, "A longitudinal study of Facebook, LinkedIn, & Twitter use," Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, pp. 2741-2750, 2012.
- [5] T. Sakaki, M. Okazaki, and Y. Matsuo, "Earthquake shakes Twitter users: real-time event detection by social sensors," Proceedings of the 19th International Conference on World Wide Web, pp. 851-860, 2010.
- [6] Y. Yorozu, M. Hirano, K. Oka, and Y. Tagawa, "Electron spectroscopy E. Kouloumpis, T. Wilson and J. D. Moore, "Twitter sentiment analysis: the good the bad and the OMG!" Proceedings of the Fifth International Conference on Weblogs and Social Media, pp. 164, 2011.
- [7] K. L. Liu, W. J. Li and M. Guo, "Emoticon smoothed language models for Twitter sentiment analysis," Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence, pp. 1678-1684. 2012.
- [8] Twitter tweet length limit, <https://dev.twitter.com/basics/counting-characters>, retrieved on 07/28/2017.
- [9] B. Pang and L. Lee, "Opinion mining and sentiment analysis," Foundations and Trends in Information Retrieval, volume 2, issue 1-2, pp. 1-135, 2008.
- [10] B. Pang and L. Lee, "Opinion mining and sentiment analysis," Foundations and Trends in Information Retrieval, volume 2, issue 1-2, pp. 1-135, 2008.
- [11] A. Esuli and F. Sebastiani, "SentiWordNet: A publicly available lexical resource for opinion mining," In Proceedings of the 5th Conference on Language Resources and Evaluation, vol. 10, pp. 2200-2204, 2006.
- [12] J. Bollen, H. Mao and X. Zeng, "Twitter mood predicts the stock market," Journal of Computational Science, vol. 2, issue 1, pp. 1-8, 2011.
- [13] I. Vegas, T. Tian, and W. Xiong, "Characterizing the 2016 U.S. presidential campaign using Twitter data," Journal of Advanced Computer Science and Applications, vol. 7, no. 10, 2016.
- [14] T. Zeitzoff, "Using social media to measure conflict dynamics," Journal of Conflict Resolution, vol. 55, issue 6, pp. 938-969, 2011.
- [15] A. Go, R. Bhayani, and L. Huang, "Twitter sentiment classification using distant supervision," Technical report, 2009.
- [16] D. Davidov, O. Tsur and A. Rappoport, "Enhanced sentiment learning using twitter hashtags and smileys," Proceedings of the 23rd International Conference on Computational Linguistics: Posters, pp. 241-249, 2010.
- [17] P. S. Dodds, K. D. Harris, I. M. Kloumann, C. A. Bliss and C. M. Danforth, "Temporal patterns of happiness and information in a global social network: Hedonometrics and twitter," PLoS ONE, vol. 6, issue 12, pp. e26752, 2011.
- [18] L. Mitchell, M. R. Frank, K. D. Harris, P. S. Dodds and C. M. Danforth, "The geography of happiness: connecting Twitter sentiment and expression, demographics, and objective characteristics of place," PLoS ONE, vol. 8, issue 5, pp. e64417, 2013.
- [19] Twitter streaming API, <https://dev.twitter.com/streaming/overview>, retrieved on 07/28/2017.
- [20] Tweeting with Location, <https://dev.twitter.com/overview/terms/geo-developer-guidelines>, retrieved on 07/28/2017.
- [21] Population estimates from United States Census Bureau, <https://www.census.gov/data/tables/2016/demo/popest/state-total.html>, retrieved on 07/28/2017.
- [22] G. F. Jenks, "The data model concept in statistical mapping", International Yearbook of Cartography, vol. 7, pp. 186-190, 1967.