

The Effect of Trending World Events on Sentiment Analysis and Relevance Intervals Using Data Analytics Software on Twitter Data

Bridgette Ravindra

Spring Valley High School, Columbia, SC 29229

Data analytics is emerging as a critical field to intelligently utilize the vast trail of data we create in our daily lives. An analysis of data trends can reveal patterns that can predict human behavior in areas such as health care, Ecommerce and consumerism, among others (Kim, 2017). The purpose of this experiment was to study the correlation between a Twitter hashtag's sentiment and its trending duration using IBM Watson Analytics. The hypothesis was that a major event associated with a more positive sentiment would trend longer than more negatively associated counterparts. The experiment relates to Hedonic adaptation, the psychological theory that states that humans will return to a relatively happy state despite a negative or positive turn of events (Halvorson, 2012). The sentiment was first analyzed on a smaller scale by randomly selecting 30 tweets within each hashtag studied and then on a larger scale using IBM Watson Analytics. For the trend analysis test, the total number of tweets for each hashtag was recorded daily. Manual sentiment analysis yielded a strong correlation of "happy" sentiment with entertainment hashtags, "sad" with natural disaster, "fearful" with health and medicine, and "neutral" with the control group #selfie. A Chi Square Test for Independence was run at $\alpha = 0.05$ on the average number of tweets for the hashtags in each category and showed a direct correlation between the category and sentiment $\chi^2(15, N = 120) = 37.731, p < 0.05$. Thus, the hypothesis was supported because the entertainment hashtags with positively associated sentiments trended longer than more serious hashtags exhibiting negative sentiments, and there was a direct correlation between the category of the tweet and its sentiment.

Introduction

Social media has become an integral component of daily life, allowing for the exchange of information and opinions across cultures and continents at a rapid rate (Kim, 2017). Among popular social media platforms such as Facebook, Instagram, Twitter and Snapchat, Twitter has firmly established itself as the hub of live, global discussion. Growing rapidly since its inception in 2006, Twitter now has over 328 million users with an estimated 80% of those users utilizing the mobile app to easily tweet what they want, when they want (Hannam, 2017). Because of the 280 character restriction on tweets, users have come to rely on syntax features, such as emoticons and hashtags, to allow their followers to better understand their mood (Bannister, 2015). Hashtags (#) are a Twitter feature that allow tweets on the same subject to be grouped and subsequently analyzed. Until recently, the aggregation, analysis, and illustration of data from social media platforms, such as Twitter, required complex programming and advanced data management skills (DeMers, 2016). However, today, there are many online services, which are readily available and user friendly to help make sense of social media data. Through the use of these data analytics programs, analysis can be performed on a large number of tweets, which when studied collectively, can contain a wealth of information that may be overlooked when observing tweets individually. Data analytics programs use linguistics to derive personality insights such as anxiety or depression, which can then be used by companies to target users for certain products (Kiser, 2017). Through analytics, it is also now possible to obtain a broader view of how people react to various situations, such as a natural disaster or a trending event in popular culture.

Research has revealed fluctuating human reactions to social and economic events worldwide and how certain days of the year can contribute to the overall mood of a large population of people (Bollen, Mao, & Pepe, 2014). For example, Thanksgiving day exhibited a sharp increase in the sentiment "vigour", which indicates physical strength and good health. Different situations or days of the year showed similarities in mood and emotional reactions. Another study focused on categorizing different tweets as being either objective or subjective, and further dividing the subjective tweets as positive or negative (Barbosa & Feng, 2010). The researchers classified the tweets through the use of syntax features, such as retweet, hashtag, reply, link, punctuation, emoticons, and upper-case letters. This is germane to this research extension because the experimenters wanted to understand a more abstract representation of tweets outside of the 280-character constraint. Similarly, another study was conducted using sentiment analysis by a three step process of tokenization, normalization, and part of speech tagging (Kouloumpis, Wilson, & Moore, 2011). From there, it was determined how effective the use of hashtags and emoticons were though the data collected from the three step sentiment analysis process. It was concluded that the part of speech analysis was not as effective as the presence of hashtags and emoticons.

One study compared the efficiency and accuracy levels of two sentiment analysis tests—Support Vectors Machine Classifiers (SVM) and the Naïve Bayes Classifier—through the use of a unigram model (Anjali, Pati, & Tripathi, 2017). After going through a stage of preprocessing, the Twitter application program interface (API), set a search query for a specific word. The overall sentiment of the tweet was determined by a unigram model that separated the adjective from the tweet. For example, in the tweet "Driving happy," the unigram model would keep only "happy." Another study took a more mathematical approach by proving that most tweets have an initial increase and then experience a slow, trailing decline as the content becomes increasingly less popular (Asur, Huberman, Szabo, & Wang, 2011). While studying the growth of tweets over time, the researchers found that the decay function $\gamma(t) = 1/t$ results in a linear increase in all trends. Similarly, another study classified tweets by using a table that presented the semantic score for a specific adjective, verb, or adverb (Kumar & Sebastian, 2012). Two functions were used in the analysis. One determined the similarity between the words and the other calculated the overall sentiment of the tweet. The results showed that this analysis was substantially effective in determining the overall sentiment.

This study focused on two measurable characteristics embedded in Twitter Streams—sentiment and trend analysis. Sentiment analysis was inferred from the total number of tweets within each sentiment, a way of determining the users' emotions using key words and syntax features, while the trend analysis, the amount of a time a hashtag is frequently used, was simply displayed on a line graph with the total number of tweets over a time period of 30 days. These analyses were conducted in a series of data analytics tests using a Twitter search query within an IBM Watson Analytics Bluemix account.

The purpose of this experiment was to compare the time period that people focus on an event perceived as consequential, such as a natural disaster, to a more relatively trivial event, such as the current popularity of a trending app. This information can help researchers gain insight into human behavior on social media and discern patterns to better understand and even predict future behavior (Aslam, 2017). Patterns of human behavior and mood can be discerned relating to core beliefs, health, political leanings, prejudices, etc., which can then be used as predictive tools in the areas of

healthcare, consumerism, and activism among others (Kim, 2017). Typically, people hyperfocus on one event and take action on it. An example of this is when Facebook users were altering their profile picture color schemes in solidarity for the Paris terrorist attack. However, this study may also show that after the initial shock, the strong focus on the event and the victims seems to fade away. This scenario relates to the psychological theory of Hedonic adaptation that states that humans will return to a relatively happy state despite a negative or positive turn of events (Halvorson, 2012).

It was hypothesized that if major events were studied through the use of IBM Watson Analytics on Twitter, hashtags with a more positively associated sentiment would trend for a longer period of time on social media than their more negatively associated counterparts. A Twitter account and IBM Watson Analytics account were created for experimentation purposes. For purposes of this experiment, trending events were classified into three main categories—natural disasters, health and medicine, and entertainment. An example of a natural disaster being studied was the Houston flooding. Health covered a range of events, such as the Ebola outbreak and the Zika virus, and entertainment covered happenings like the viral popularity of the app, *Pokémon GO*, or the movie, *Finding Dory*. Some of the most popular hashtags within each category were selected. Since the category was altered throughout the trials of the experiment, the major event studied was considered the independent variable, and the dependent variable was the total number of tweets within each sentiment and the total number of tweets containing each hashtag. The control group consisted of 30 trials with no major event associated with it. In this case, the hashtag, #selfie, was used.

Methods

An IBM Bluemix and Twitter account were created prior to experimentation. Three categories of tweets were developed—natural disasters, health and medicine, and entertainment, and three popular hashtags from each category were chosen. The natural disaster tweets included the hashtags, #houstonflood, #hurricanemattthew, and #scfloods. The health and medicine tweets encompassed #zika, #opioidepidemic, and #ebola. The last category, entertainment, included tweets with #pokemongo, #findingdory, and #rio2016. The control group were tweets with no specific event associated with it. For purposes of this experiment, the hashtag, #selfie, was used.

The sentiment analysis test was conducted first. It was first carried out on a smaller, more precise level, and then compared to a broader sentiment analysis result set. The data was refined in the IBM Watson account, so all of the tweets were displayed from the date it occurred and for 10 days following. For example, the movie release date was used for #findingdory and continued for 10 days. A random number generator was used to choose 30 numbers, and the numbers were matched to the tweets in the data set. From there, data was recorded in a spreadsheet, including the category, hashtag, date parameters, number of tweets available, location and username, and the actual tweet. Each tweet was manually categorized into as specific sentiment. For purposes of this experiment, “happy” was represented by yellow, “angry” with red, “sad” with blue, “fearful” with green, “affectionate” with purple, and “neutral” with orange. Each cell in the spreadsheet was color coded appropriately, and a bar graph was created displaying the sentiment and the number of tweets within each sentiment for each hashtag. The sentiment analysis was then conducted on a larger scale using the IBM Watson account. A sentiment bar graph was created to show which sentiments were highest for each hashtag using all of the tweets from the 10 day period.

Next, the trend analysis test was conducted. The data was scraped from the Twitter website through the use of Python code written in DSX. This pulled data from Twitter using their publicized method of scrolling through tweets. The data was then filtered, so it displayed only the date and the number of tweets containing the hashtag on that date. A line graph was created to show the major peaks and falls of the hashtag over a period of 30 days.

Results

Four categories were used in experimentation—entertainment, natural disaster, health and medicine, and the control. Within each Twitter category, thirty randomly selected tweets were analyzed from the larger data set generated by Watson Analytics. For each sentiment within the different categories, an average was calculated from the three hashtags in that category (Table 1). The highest number of “happy” tweets were associated with the entertainment hashtags with an average of 13 tweets for each hashtag. The health and medicine category had the highest number of “angry” tweets with an average of 5.667 tweets per hashtag. The “sad” sentiment was most closely related to the natural disasters category with an average of 7.333 tweets for each hashtag. Health and medicine had the highest number of “fearful” tweets with an average of 6.333 tweets per hashtag. The highest number of “affectionate” tweets was associated with the natural disasters category with an average of 4 tweets per hashtags. The control had the most “neutral” tweets with a total of 17 tweets for #selfie.

Table 2 displays the descriptive statistics using the data from Table 1. Results show the mean to be highest for the health and medicine category with an average of 5.17, while the other categories had a mean of 5. The greatest data variation can be seen in the control category with the highest SE mean of 3.080 and a standard deviation of 7.540. The next greatest variation was in the entertainment, then health and medicine, and finally, natural disaster with the lowest SE mean of 0.807 and a standard deviation of 1.978.

Figure 2 displays data collected from the manual sentiment analysis test. Each bar represents the total number of tweets per sentiment, and the different colors represent the categories’ average number of tweets per sentiment. “Neutral” and “happy” were the highest sentiments, with the largest percentage of their bars being the control and the entertainment hashtags. “Affectionate” was the lowest sentiment with the largest portion of the tweets being associated with a natural disaster hashtag.

After the manual sentiment analysis test, the study was conducted on a larger scale using IBM Watson Analytics. The sentiment analysis bar graphs produced using Watson are displayed in Figure 3. The highest sentiment of each hashtag was determined to be “neutral.” Overall, the entertainment and natural disaster hashtags had the second highest sentiment of “positive.” The health and medicine hashtags had the second highest number of “negative” tweets.

Figure 4 displays the total number of tweets per day for each hashtag. Figure 4.1 shows the trending line for each entertainment hashtag. These hashtags had the highest peaks and were still trending after a period of thirty days. The total number of tweets for each natural disaster hashtag is displayed in Figure 4.2. These hashtags were used less frequently and were diminishing in popularity by the end of the thirty day period. Figure 4.3 shows the total number of tweets for the health and medicine hashtags. These hashtags also trended throughout the thirty days, and were still fairly popular towards the end. Trend analysis for the control, #selfie, can be seen in Figure 4.4. The total number of tweets remained relatively consistent compared to the other experimental hashtags.

Table 3 displays the data from the Chi-Square Test for Independence between the category of the tweet and the overall sentiment of the tweet. Using the average number of tweets per sentiment from the three hashtags in each category, this statistical test was run with a null hypothesis that

sentiment and the category of the tweet were not related. The alternative hypothesis was that the category and sentiment were related. The Chi Square Test showed that there was a direct correlation between the category and sentiment $X^2(15, N = 120) = 37.731, p < 0.05$

Discussion

The purpose of this experiment was to explore how the sentiment associated with an event impacts how people react to it on social media, specifically Twitter. Through social media analytics, researchers can gain insights into human behavior that can be leveraged to predict trends in areas such as Ecommerce and consumerism, health and wellness, and science and technology. This research focused on the emerging field of sentiment analysis to assess how human behavior is impacted by current events. It was hypothesized that if major events were studied on social media, hashtags associated with a positive sentiment would trend longer than more consequential events associated with a negative sentiment. This hypothesis was supported: hashtags associated with positive sentiments in the entertainment group corresponded with longer trending times extending beyond the 30 day trial period. This also substantiates the theory of Hedonic adaptation and how people tend to dismiss negative events and return to a stable state of happiness.

A Chi Square test for independence was run at $\alpha = 0.05$ to determine whether a correlation existed between the average number of tweets for each hashtag within the different categories and the overall sentiment of the tweet. Results showed that there was a direct correlation between the category and sentiment with $X^2(15, N = 120) = 37.731, p < 0.05$. This means that the most positively associated category, entertainment, had the highest average number of “happy” tweets, while the negatively associated category, natural disasters, had the highest average number of “sad” tweets.

The data suggest that events and hashtags associated with more positive sentiments trend for a longer period of time than those with negatively associated hashtags. Findings of this study indicate that mood can significantly impact human behavior on social media. Researchers have determined that the way the human mind reacts to social and economic events also depends on the day of the year and general population mood (Bollen, Mao, & Pepe, 2014). For example, students and professors may feel more positive during the summer months when school is not in session. This directly relates to the research project because it shows how certain sentiments are prominent during various times of the year. If a major natural disaster occurred over the summer, people may feel more relaxed about the event due to their current surroundings. Barbosa and Feng (2010) analyzed sentiment by categorizing tweets as objective or subjective and then into positive or negative. Using a similar syntax analysis approach as this research project, such as retweet, hashtag, punctuation, emoticons, and upper-case letters, tweets were classified into varying sentiments. Similarly, another study ran sentiment analysis through a three step process approach—tokenization, normalization, and part of speech tagging (Kouloumpis, Wilson, & Moore, 2011). This was compared to another test analyzing sentiment through the use of syntax features. Results showed that the part of speech analysis was not as effective as searching for key words, hashtags, and emoticons. By analyzing these trends with sentiment analysis, many useful inferences can be made in the study of human behavior.

Other algorithms and mathematical equations can prove that over time, tweets will have an initial increase and then trail off. A linear increase in trends using the decay function $\gamma(t) = 1/t$ is demonstrated for various hashtags (Asur, Huberman, Szabo, & Wang, 2011). Although this linear graph was not demonstrated in this study, the data showed the trend lines to be relatively consistent for each category of tweet. Kumar and Sebastian (2012) classified tweets using a table that displayed the semantic score for a specific adjective, verb, or adverb. Results showed that using two functions—one that determined the similarity between the words and one that calculates the overall sentiment of tweet—is a very effective way to determine tweet sentiment. This study again shows that through the use of key words and syntax features, the overall sentiment can be determined.

Results corroborated that the overall sentiment of a hashtag is directly correlated to the trending time. The entertainment hashtags were most often associated with the “happy” sentiment, and they continued to trend even after the 30 day period. The natural disaster hashtags exhibiting the mostly “sad” sentiment trended for the shortest amount of time. The health and medicine hashtags yielded the second longest trending period of time and were associated with the “fearful” sentiment, and the control group trended for a comparatively consistent time and had the most “neutral” tweets.

There are a number of sources of error in this experiment. The most significant one was the constraints imposed by Twitter on accessing private user accounts. However, it should be noted that public accounts make up a substantial percentage of overall Twitter accounts and should reflect an accurate sample of the dataset as a whole. In addition, certain adjustments to the methods could be made to improve this research. Events in closer proximity to one another could be chosen yielding greater consistency in factors such as the number of active twitter users. During the manual sentiment analysis test, the researcher categorized the tweet into different sentiments. This, however, may be subjective since the overall sentiment of a tweet is opinionated. Instead, using a team of researchers that could all give their input on the overall sentiment would ensure more accurate results. The time parameters may have also resulted in imprecise results. The sentiment analysis was carried out over a period of 10 days, but the overall sentiment may shift over an extended period of time as public opinion wavers. For example, the natural disaster tweets may have become more “happy” over time as relief efforts begin. The trend test was capped after a period of 30 days. This, however, may have been inaccurate as some of the hashtags may have dropped off or peaked outside of these time parameters. To improve this, research could be carried out over an extended period of time. Unfortunately, data size constraints prohibited such time parameters in this experiment.

Future research could involve studying hashtag trends on a larger scale over a longer period of time if file size limitations could be resolved. This would allow for a more accurate representation of the major peaks and falls of the various tags. Different categories of tweets could also be examined, such as politics, to see how sentiment analysis can be used as a predictor for elections. In addition, the study could focus on certain geographical areas of the world to see how the popularity of certain events in the United States compares to other countries. Other analytics applications, aside from sentiment analysis on social media, could include researching consumer purchasing patterns to improve product marketing or pricing.

Acknowledgements

I would like to thank Mrs. Michelle Spigner for her guidance and assistance throughout the course of developing this research project, experimentation, and statistical analysis. I would like to acknowledge Vishal Chahal, a data scientist at IBM Watson, for advice and guidance. I would also like to offer my sincere gratitude towards my peers who gave me their opinions and suggestions on how to improve my work in its entirety. Lastly, I would like to thank my parents and brothers for their assistance and support.

Notes and References

1. Halvorson, H. G. 2012. *Psychology Today*. New York (NY): Sussex Publishers, LLC. [cited 18 October 2017]. Available from <https://www.psychologytoday.com/blog/the-science-success/201208/How-keep-happiness-fading>
2. Hannam, S. 2017. *Reaching Today's Tech Savvy On-The-Go-Families*. [cited 19 October 2017]. Available from http://www.niagararegionpta.org/files/6015/0402/3885/NRPTA_Social_Media_-_Fall_2017.pdf
3. Bannister, K. 2015. *Brandwatch*. Brighton (UK): Brandwatch. [cited 18 October 2017]. Available from <https://www.brandwatch.com/blog/understanding-sentiment-analysis/>
4. DeMers, J. 2016. *Forbes*. Jersey City (NJ): Forbes Media, LLC. [cited 18 October 2017]. Available from <https://www.forbes.com/sites/jaysondemers/2016/02/09/how-long-will-todays-social-media-platforms-last/#13ced0b71b72>
5. Kiser, M. 2017. *Algorithmia*. [cited 19 October 2017]. Available from <https://blog.algorithmia.com/introduction-sentiment-analysis-algorithms/>
6. Bollen, J., Mao, H., & Pepe, A. 2014. Modeling public mood and emotion: Twitter sentiment and socio-economic phenomena. *The Association for the Advancement of Artificial Intelligence*. Available from <https://www.aaai.org/ocs/index.php/ICWSM/ICWSM11/paper/viewFile/2826/3237>
7. Barbosa, L. & Feng, J. 2010. Robust sentiment detection on Twitter from biased and noisy data. *The Association for Computational Linguistics*. Available from <http://www.aclweb.org/anthology/C10-2005>
8. Kouloumpis, E., Wilson, T., & Moore, J. 2011. Twitter sentiment analysis: The good, the bad, and the OMG! *The Fifth International AAAI Conference on Weblogs and Social Media*. Available from <https://www.aaai.org/ocs/index.php/ICWSM/ICWSM11/paper/view/2857/3251>
9. Anjali, M., Pati, P., & Tripathi, A. 2017. Twitter sentiment analysis of current affairs. *International Journal on Recent and Innovation Trends in Computing and Communication*, 5(5). Available from http://www.ijritcc.org/download/browse/Volume_5_Issues/May_17_Volume_5_Issue_5/1496050566_29-05-2017.pdf
10. Asur, S., Huberman, B. A., Szabo, G., & Wang, C. 2011. Trends in social media: Persistence and decay. *HP Labs*. Available from http://www.hpl.hp.com/research/scl/papers/trends/trends_web.pdf
11. Kumar, A. & Sebastian, T. M. 2012. Sentiment analysis on Twitter. *International Journal of Computer Science Issues*, 9(4). Available from <https://www.ijcsi.org/papers/IJCSI-9-4-3-372-378.pdf>
12. Aslam, S. 2017. *Omnicores*. [cited 19 October 2017] Available from <https://www.omnicoreagency.com/twitter-statistics/>
13. Kim, L. 2017. *HubSpot*. [cited 2 November 2017]. Available from <https://blog.hubspot.com/marketing/twitter-analytics>

APPENDIX A

Figure 1: Experimental Design Diagram

Title: The effect of trending world events on sentiment analysis and relevance intervals using data analytics software on Twitter data									
Hypothesis: If major events are studied through the use of IBM Watson Analytics for social media, hashtags with a more positively associated sentiment analysis will trend for a longer period of time on social media than their more serious counterparts that have a more negatively associated sentiment analysis.									
Independent Variable: the major world event									
Levels of IV	Natural Disaster Tweets			Health and Medicine Tweets			Entertainment Tweets		
	#houston flood	#hurricane matthew	#scfloods	#zika	#opioidepidemic	#ebola	#pokemongo	#finding dory	#rio2016
Number of Repeated Trials (number of tweets)	30	30	30	30	30	30	30	30	30
Dependent Variable: the total number of tweets within each sentiment and the total number of tweets containing each hashtag									
Constants: data analytics software used, social media site used, and amount of time per trial for sentiment and trend analysis									

Table 1. Average number of tweets for the hashtags within each category of tweet

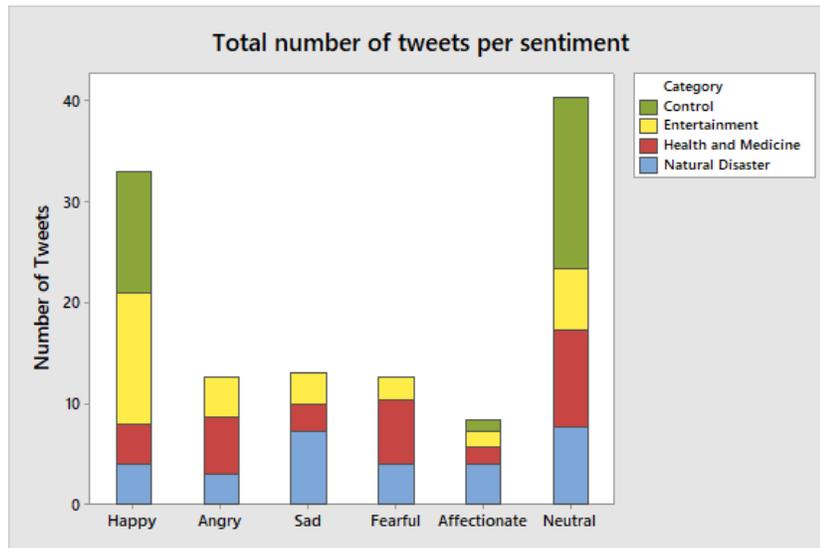
	Happy	Angry	Sad	Fearful	Affectionate	Neutral	Total
Natural Disasters	4	3	7.333	4	4	7.667	30
Health and Medicine	4	5.667	2.667	6.333	1.667	9.667	30
Entertainment	13	4	3	2.333	1.667	6	30
Control	12	0	0	0	1	17	30
Total	33	12.667	13	12.666	8.334	40.334	120

Table 1 displays the average number of tweets for each hashtag associated with each sentiment. Entertainment had the highest number of “happy” tweets, while health and medicine had the greatest average for the “angry” sentiment. Natural disasters had the most “sad” tweets, and health and medicine had the greatest number of “fearful” tweets. “Affectionate” had the highest number of tweets associated with natural disasters, while the control had the most “neutral” tweets.

Table 2. Descriptive statistics for the average number of tweets for the hashtags in each category of tweet for every sentiment (Table 1)

Category of Tweet	Mean	SE Mean	Standard Deviation	Range	Min	Q1	Q3	Max
Natural Disaster	5.00	0.807	1.978	4.667	3	3.750	7.417	7.667
Health and Medicine	5.17	1.170	2.860	8.000	2	2.750	7.000	10.000
Entertainment	5.00	1.710	4.200	11.000	2	2.000	7.750	12.000
Control	5.00	3.080	7.540	17.000	0	0.000	13.250	17.000

The greatest variation in the data starts with the control category, then entertainment, followed by health and medicine, and ending with natural disaster, showing the lowest variation. The health and medicine category had the greatest mean of 5.17, while the other categories had an average of 5.

Figure 2. Sentiment analysis with average number of tweets for each category per sentiment

The average number of tweets per hashtag was calculated for each hashtag and displayed using a different color within each bar, representing the total number of tweets per sentiment. “Happy” and “neutral” were the highest sentiments, most popular in the control and entertainment tweets. “Affectionate” was the lowest sentiment, and this sentiment was most closely associated with the natural disaster hashtags.

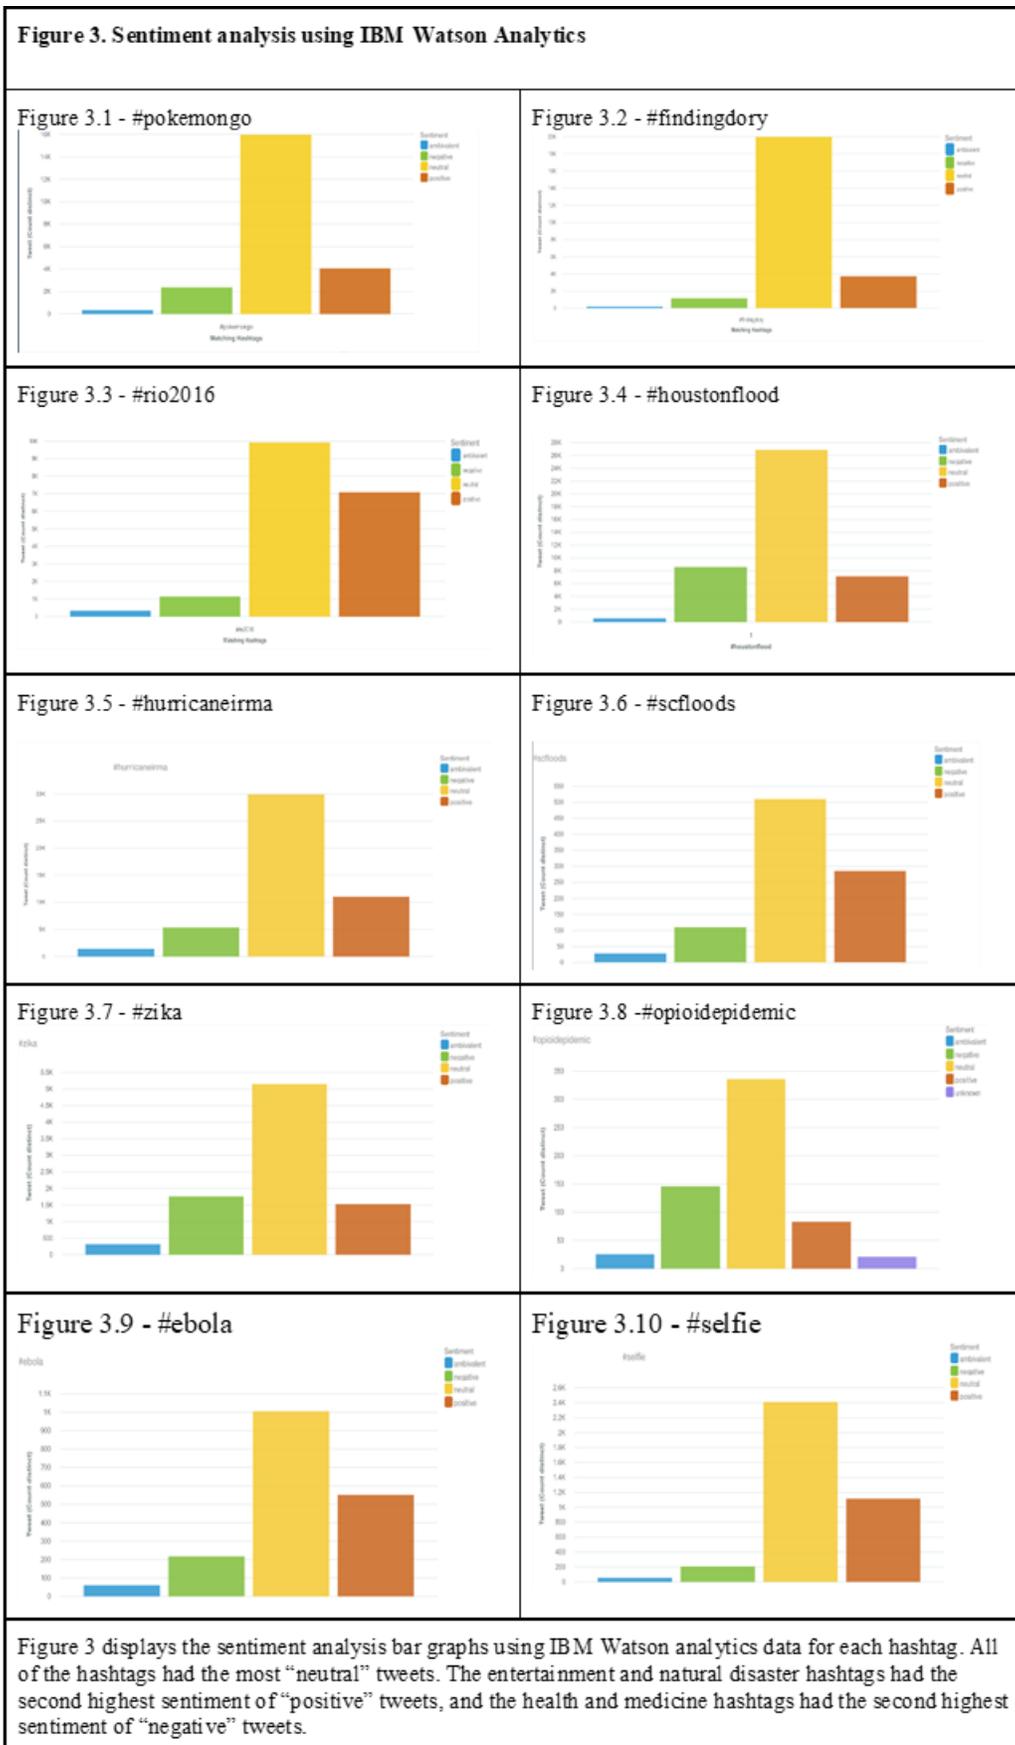


Figure 4. Trend analysis for each category of hashtag

Figure 4.1

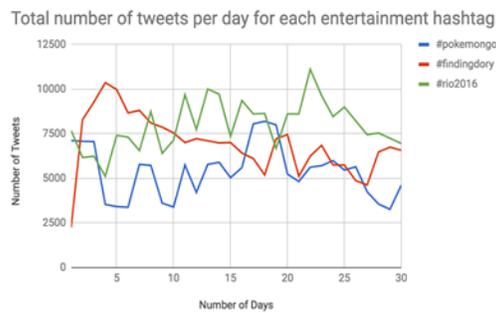


Figure 4.2

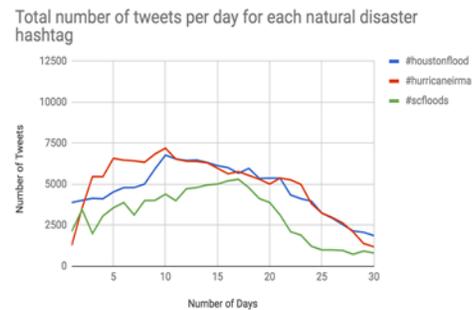


Figure 4.3

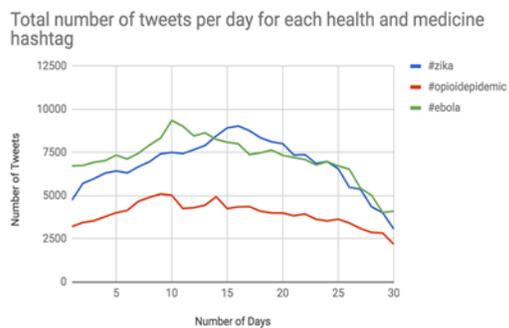


Figure 4.4

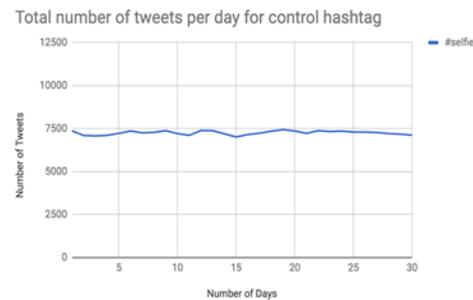


Figure 4.1 displays the total number of tweets each day containing the entertainment hashtags. These hashtags were used most frequently and were still trending after thirty days. The total number of tweets with natural disaster hashtags can be seen in Figure 4.2, and Figure 4.3 shows health and medicine data. The health and medicine hashtags trended for the second longest time, and the the natural disaster hashtags trended for the least amount of time. The control hashtag (Figure 4.4) remained relatively consistent.

Table 3. Chi-Square test for independence results for the average number of tweets per sentiment from the three hashtags in each category over a period of 10 days (Table 1)

	Chi-Square	DF	P-value
Pearson	37.731	15	0.001
Likelihood Ratio	44.880	15	0.000

Table 3 shows the data observed from the Chi-Square test for independence run for the average number of tweets per sentiment from the hashtags in each category. Run at alpha = 0.05, the Chi Square Test showed that there was a direct correlation between the category and sentiment $X^2(15, N = 120) = 37.731, p < 0.05$.