EMOTIONS AND ELECTIONS: EVIDENCE FROM HURRICANE SANDY

PRABHAT BARNWAL, JAMES RISING

October 26, 2014

Abstract. Social media offers a detailed window into the responses of a population to disasters. Measuring such response is essential to understand how extreme events indirectly affect real world events. This paper uses time and space variation induced by hurricane Sandy in Oct-Nov 2012 to measure how natural disasters influence the emotions expressed by people on the social media. We employ micro-blogging data from about 2 million geolocated entries in Twitter and develop a methodology to estimate emotions from tweets. Our preliminary results show proportional and symmetric effects driven by the variation in magnitude and timing of weather events, particularly wind speed and precipitation. Next, we extend the analysis to realtime election data, using the prediction market prices from online platform InTrade and pre-poll surveys, to suggest how these emotional reactions might have influenced the 2012 presidential election outcome.

1. Literature

Twitter, a social media platform, can act as a powerful window into social dynamics, due to the public availability of tweets and their real-time use (Kwak et al., 2010). A vibrant literature is emerging on the potential for social media in disaster situations. Recent work has explored Twitter adoption (Hughes and Palen, 2009), information exchange (Mendoza et al., 2010, Heverin and Zach, 2010, Muralidharan et al., 2011), situational awareness (Vieweg et al., 2010, Acar and Muraki, 2011). Emotion has also been a key area within Twitter research. The primary interest in many emotion papers is to study how users influence each other (e.g., Quercia et al., 2011, Kivran-Swaine and Naaman, 2011, Jansen et al., 2009).

Few papers, however, have focused on emotions during disasters as expressed in social media. Qu et al. (2009) shows that users do use social media to express their emotions about disasters, and Balog et al. (2006) shows that these emotional sentiments can be mapped to events in news.

This paper also explores the connections between emotions around Sandy and the subsequent election. Tumasjan et al. (2010) previously showed that simple mentions of political parties correlate with those party's future performance.

2. Collection Method

The results in this working paper come from an analysis of 11251 tweets collected in real-time. Future analysis will encompass a collection of 2 million tweets, representing every geolocated on the Northeastern seaboard between Delaware and southern Connecticut, from October 20 to November 5.

Twitter data were collected by searching for 100 tweets posted within 10 km of 40.714353° N, -74.005973° W (near New York City Hall), every hour between October 24, 2012, 2 am GMT (Oct. 23, 10 pm, NYC local time) and November 1, 2012, 3 pm GMT (Nov. 11, 11 am, NYC local time). During this period, 20436 unique tweets were collected.

Of these, 314 tweets were posted with latitude and longitude information. Across the remaining tweets, there are 6844 different location labels. 11251 tweets were categorized into ten regions. These are shown in figure 1.

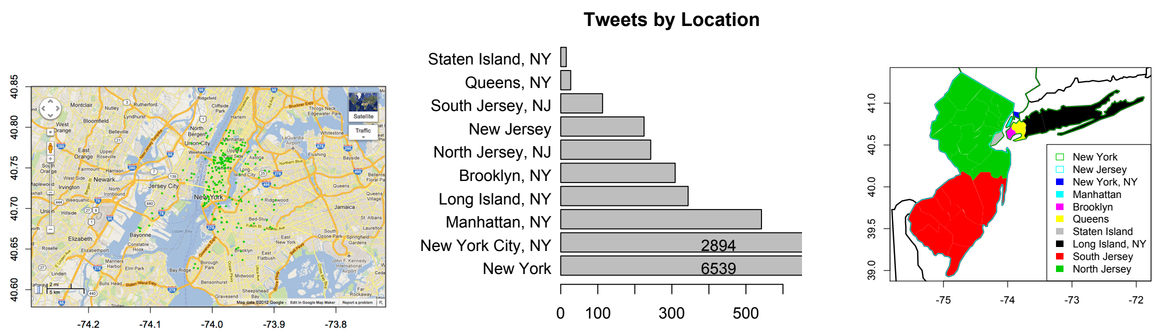

Geocoded tweet locations All determinable locations Region locations

FIGURE 1.

3. Sentiment Analysis

The emotional content for each tweet is estimated on eight non-orthogonal axes. We combine the Affective norms for English words (ANEW) dataset (Bradley and Lang, 1999) with finer-grained emotions using the models from Stevenson et al. (2007). Both of these collections are based on single-word emotional responses, as reported by experimental subjects.

For determining emotional content, each tweet is treated as a set of words. Valence, arousal, and dominance values are identified for words in the ANEW dataset, or for their stems using the Porter stemmer algorithm (Porter, 1980). The full tweet's values for valence, arousal, and dominance are then the average across all identified words or stems.

$$v_i = \frac{1}{m_i} \sum_{\text{word} \in \text{identified}} V(word)$$

where $V(\cdot)$ provides a 3 element vector of ANEW emotion values for each word, m_i is the number of identified words in tweet i, and v_i is a 3 element vector of ANEW emotion values for the entire tweet

Stevenson et al. (2007) provides models of the three ANEW emotional ratings, based on participant's ratings of words for happiness, anger, sadness, fear, and disgust. We use a minimum length estimation to determine the five emotional ratings from the three provided by ANEW, as follows. The original models in Stevenson et al. can be expressed as,

$$v = Ge$$

where v is a 3xN matrix of ANEW emotion estimates, e is a 5xN matrix of (unknown) Stevenson emotions, and G is a 3x5 matrix of linear model coefficients. The minimum length solution is,

$$e = G^\top \left(GG^\top\right)^{-1} v$$

This provides the mean estimate. We also produce a clipped Gaussian approximation to the standard error distribution, by treating each element of v, e, and G as a random variable with a Gaussian distribution clipped between 0 and 1. After the arithmetic operations described above are applied to these distributions, the final emotional value estimates are the mean of the resulting clipped Gaussian distributions.

We also impute the emotional content of unknown words recursively, with a simplified EM algorithm. After first estimating the emotional content of all tweets as above, we treat the all unknown words as having a distribution of emotional values from the estimated values for the tweets they are represented within. This expands the collection of words with emotional values. The process is repeated , with the emotional values for both unknown words and previously estimated words being updated to reflect the new estimates.

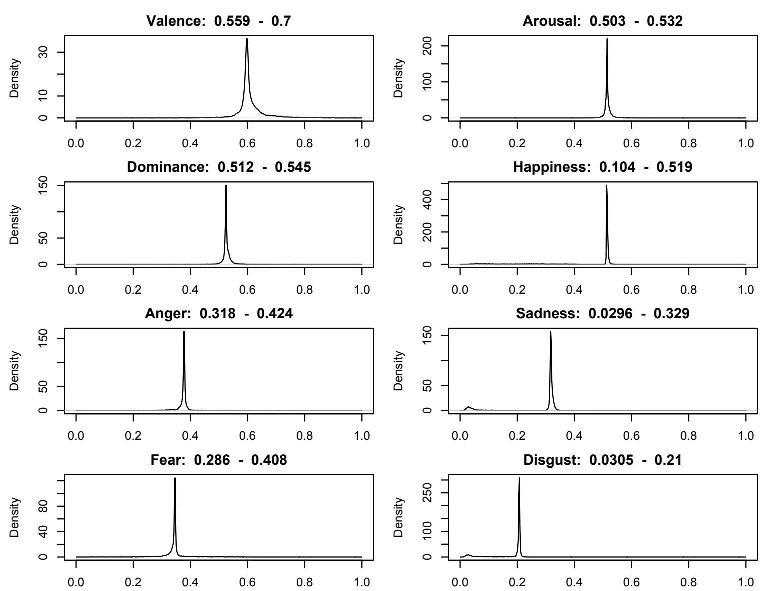

Chat symbol	Valence	Arousal	Dominance
`:)`	0.6167 (75%)	0.5153 (63%)	0.5277 (70%)
`:(`	0.5896 (19%)	0.5126 (27%)	0.5229 (27%)
`lol`	0.6042 (62%)	0.5150 (60%)	0.5258 (62%)
`wtf`	0.6026 (59%)	0.5168 (72%)	0.5265 (66%)

FIGURE 2. **Top:** Distributions for each of the emotions, with 90% confidence intervals. **Bottom:** Average values of valence, arousal, and dominance for tweets containing emoticons and chat acronyms.

The result of these operations presses the emotion vales into narrow bands for three reasons. First, the mean emotion from all of the words in a tweet follows the central limit theroem. Second, the distributions are clipped between 0 and 1, so that the mean of this distribution is held away from those extremes. This effect becomes particularly strong for the derived emotions, since the operations result in wide uncertainties, which when clipped produce means near .5. Third, imputed values tend to be near .5, since they reflect the average of many different tweets, which produces a further influence toward this middle value.

The sentiment analysis does not account for emoticons or chat acronyms. We use these as a check, as shown in figure 2. Example tweets, representing the 8 emotional poles of ANEW, are shown in table 1.

	Low valence	**High valence**
Low arousal, Low dominance	"stOMACH HURTS"	"My pillow has my heart."
Low arousal, High dominance	"I hate this 4 girls sharing one bathroom nonsense."	"Alive Thankful Blessed"
High arousal, Low dominance	"This movie is terrible."	"3..2..1, Here comes the hurricane baby"
High arousal, High dominance	"Rage is appropriate and powerful and moving."	"My birthday party was amazing. Loved it."

Table 1. Example tweets, representing the 8 emotional poles in ANEW.

4. Weather

We employ hourly recorded weather from weather stations in New York city and surrounding areas. The spatial resolution is at the borough level. Main parameters are precipitation, wind-speed, temperature, gust-speed, visibility,humidity etc. Weather data employed is from 24 Oct to 1 Nov. On Monday, 29 October at 9 PM, hurricane Sandy made a landfall in Atlantic city, NJ.

5. Analysis

We first present New York city level analysis of tends in various key sentiments around Sandy's arrival time. Because of limited geo-location identification in majority of tweets, we find it difficult to run borough level analysis.

The following graphs show mean sentiments (valence, arousal, dominance, happiness, anger, sadness, fear and disgust) of all the tweets recorded in each hour. A local polynomial of relevant weather variable is shown below. Similar polynomial for six main weather variables (wind speed, gust speed, precipitation, temperature, humidity and visibility) are also shown in next graph. It is clear that more of the sentiments respond to hurricane weather, particularly wind speed and precipitation. A joint plot showing movement in three emotions (happiness, fear and disgust) along with two weather variables (wind speed and precipitation) is shown in the next graph.

The results show that valence, dominance, happiness, anger, sadness, and disgust decrease with the advent of Sandy, while arousal and fear increase. Both happiness and sadness can simultaneously decrease, due to their similar dependence on arousal and dominance. However, the emotions do not peak simultaneously. Fear peaks just before landfall, and dominance falls to its lowest point just afterwards. Anger and arousal maintain their extreme values for the longest period.

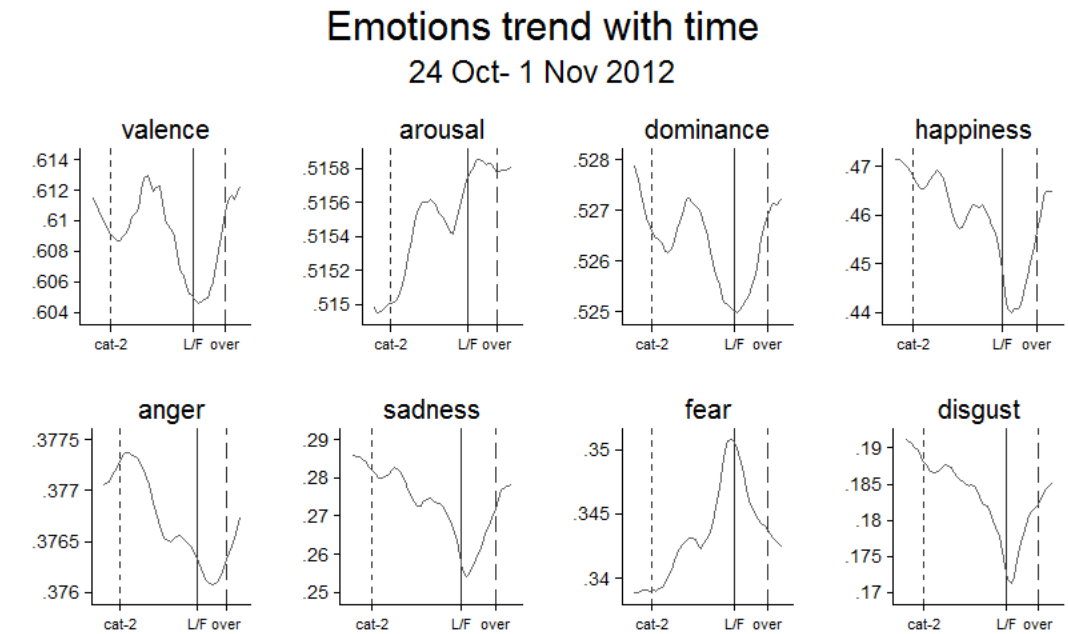

Above plots show trend in eight emotions synthesized from 18,618 tweets in New York city from Oct 24 to Nov 1 2012.
Y- axis: hourly average of respective synthesized emotions values from each tweet.
X- axis shows hours starting from Wed Oct 24 2012, 1 AM EDT
Reference lines show the timeline:
(1) Cat 2: Sandy declared Category 2 on Wed 24 Oct 2012 at 11 AM EDT
(2) L/F: Landfall in Atlantic City on Mon Oct 29 2012 at 8 PM EDT
(3) Over: Sandy completely disappeared from New York area on Wed Oct 31 2012 by noon EDT)

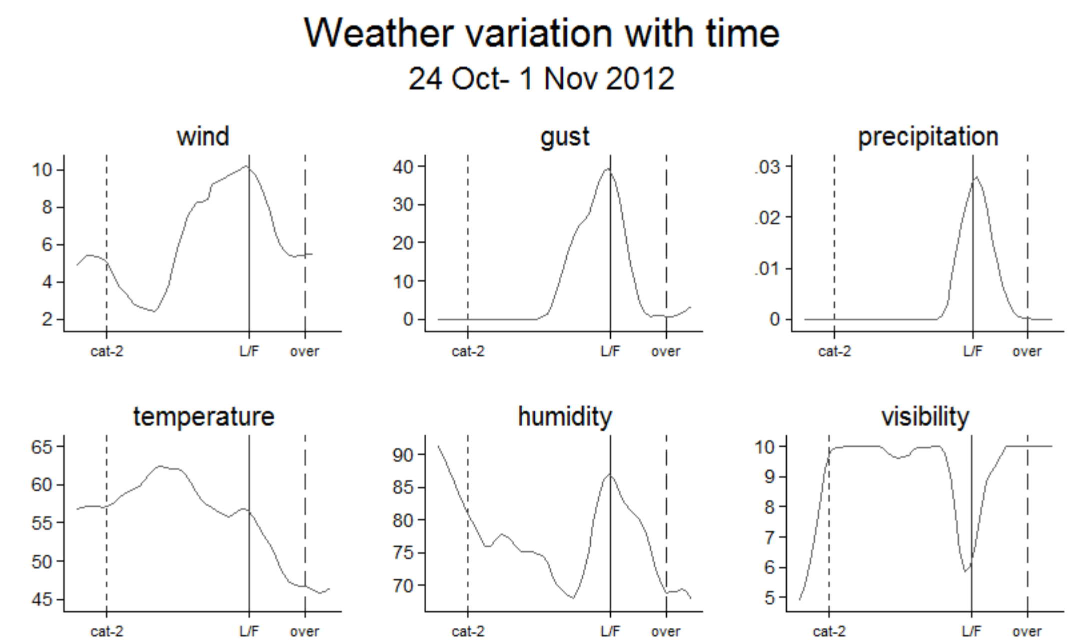

Above plots show variation in six major weather variables in New York city.
Y- axis: hourly values of respective weather variable as monitored in New York city
X- axis shows hours starting from Wed Oct 24 2012, 1 AM EDT
Reference lines show the timeline:
(1) Cat 2: Sandy declared Category 2 on Wed 24 Oct 2012 at 11 AM EDT
(2) L/F: Landfall in Atlantic City on Mon Oct 29 2012 at 8 PM EDT
(3) Over: Sandy completely disappeared from New York area on Wed Oct 31 2012 by noon EDT)

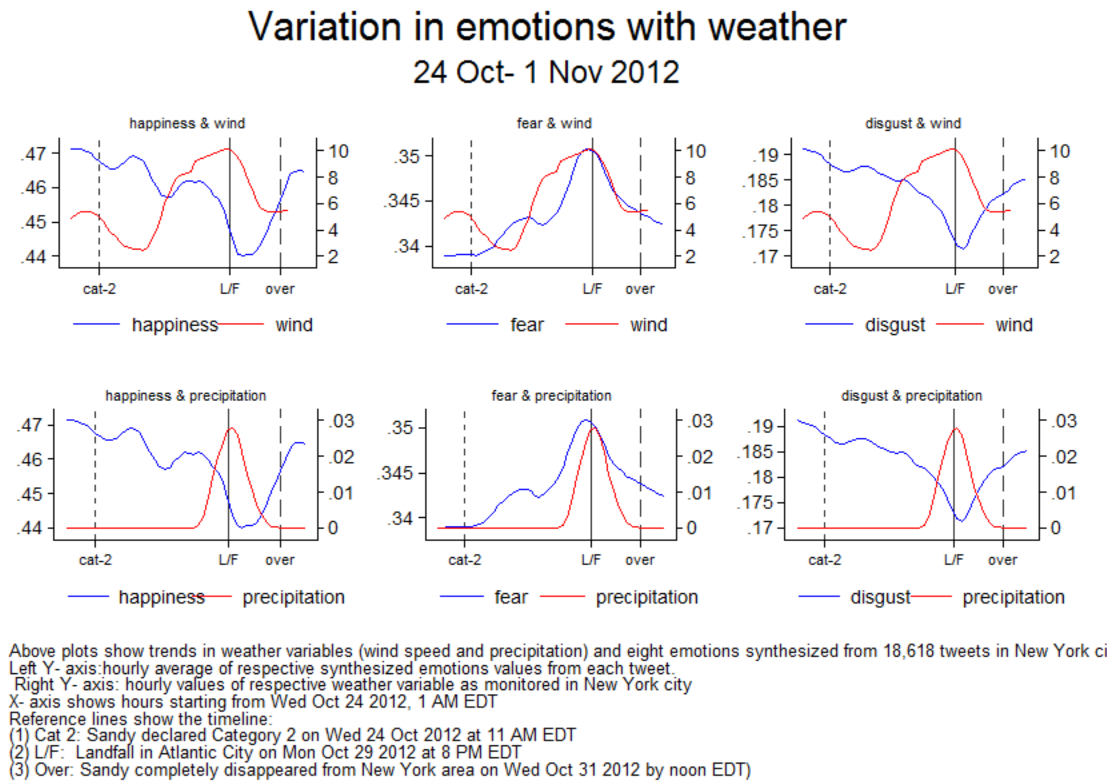

6. Summary Statistics

We collected a larger number of tweets through June to construct averages for the eight emotions. The anomalies between this average and the tweet emotions before and after Sandy (according to two dates) are shown in table 2.

	N	Valence	Arousal	Domin.	Happi.	Anger	Sadness	Fear	Disgust
October through June	539517	0.6054	0.5126	0.5278	0.4692	0.3811	0.2814	0.3387	0.1901
Before Mon, 29, anomaly	11827	0.0049	0.0027	-0.0011	-0.0043	-0.0036	-0.0026	0.0028	-0.0036
After Mon, 29, anomaly	8563	0.0021	0.0032	-0.0017	-0.0156	-0.0050	-0.0139	0.0076	-0.0105
Before Sat, 27, anomaly	7178	0.0047	0.0026	-0.0010	-0.0008	-0.0039	0.0011	0.0016	-0.0016
After Sat, 27, anomaly	13212	0.0032	0.0031	-0.0016	-0.0136	-0.0043	-0.0119	0.0066	-0.0091

Table 2. Average emotional anomalies before and after Sandy across entire sample.

While most users occur only once in our dataset, about 10% tweet twice or more. Considering only these users, we estimate the how their average emotions changed before and after the event, using two dates, in table 3.

Although strong correlations exist between weather variables and average emotions, borough-level weather is very poorly correlated with individual tweet emotions due to the high level of noise. In addition, weather variables are highly correlated with eachother. To explore the possible effects of thresholds, we search for the logistic transformations of the weather variables that have the greatest explanatory effect. Then we build linear models for each emotion, using all of these transformed weather variables. The results are in table .

Mon, 29 Oct 2012 00:01 - N = 338.

	Differences	Std. Errs	P-value	
valence	-0.008	0.004	0.006	**
arousal	0.000	0.001	0.918	
dominance	-0.002	0.001	0.045	*
happiness	-0.009	0.008	0.111	
anger	-0.002	0.003	0.378	
sadness	-0.008	0.006	0.078	.
fear	0.002	0.003	0.480	
disgust	-0.004	0.003	0.112	

Sat, 27 Oct 2012 00:59 - N = 517.

	Differences	Std. Errs	P-value	
valence	-0.008	0.003	0.002	**
arousal	-0.001	0.001	0.110	
dominance	-0.001	0.001	0.065	.
happiness	-0.009	0.007	0.050	.
anger	0.001	0.002	0.390	
sadness	-0.008	0.005	0.025	*
fear	0.002	0.002	0.151	
disgust	-0.003	0.003	0.202	

Table 3. Emotional anomalies between tweets from the same users, before and after Sandy.

	Valence	Arousal	Domin.	Happi.	Anger	Sadness	Fear	Disgust
(Intercept)	0.63^{***}	0.52^{***}	0.55^{***}	0.43^{***}	0.45^{***}	0.28^{***}	0.36^{***}	0.22^{***}
	(0.04)	(0.01)	(0.01)	(0.11)	(0.04)	(0.08)	(0.04)	(0.04)
Temperature	0.00	0.00	0.01	−0.05	0.03^{***}	−0.04	0.01	0.00
	(0.01)	(0.00)	(0.00)	(0.03)	(0.01)	(0.03)	(0.01)	(0.01)
Dew point	0.00^{***}	0.00	0.00	0.00	0.00^{**}	0.00	0.00	0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Humidity	−0.02	0.00	0.00	−0.05	0.02^{*}	−0.02	-0.03^{*}	0.02
	(0.02)	(0.00)	(0.00)	(0.04)	(0.01)	(0.03)	(0.01)	(0.02)
Sea-level pressure	0.04	-0.02^{**}	0.01	0.32^{***}	-0.07^{**}	0.26^{***}	-0.10^{***}	0.10^{***}
	(0.03)	(0.01)	(0.01)	(0.08)	(0.03)	(0.06)	(0.03)	(0.03)
Visibility	0.00	0.00	-0.01^{***}	−0.01	-0.02^{**}	−0.01	0.01	-0.02^{**}
	(0.01)	(0.00)	(0.00)	(0.02)	(0.01)	(0.02)	(0.01)	(0.01)
Wind speed	0.00	0.00	0.00	0.02^{**}	0.00	0.01^{**}	0.00	0.01^{**}
	(0.00)	(0.00)	(0.00)	(0.01)	(0.00)	(0.01)	(0.00)	(0.00)
Wind gusts	0.05	0.01	0.00	−0.09	−0.04	-0.11^{*}	0.01	-0.11^{***}
	(0.03)	(0.01)	(0.01)	(0.08)	(0.03)	(0.06)	(0.03)	(0.03)
Precipitation	−0.14	0.00	-0.06^{**}	−0.27	−0.09	-0.31^{*}	0.16^{**}	-0.19^{**}
	(0.09)	(0.03)	(0.03)	(0.23)	(0.08)	(0.18)	(0.08)	(0.10)
R^2	0.00	0.00	0.00	0.00	0.00	0.01	0.00	0.01
Adj. R^2	0.00	0.00	0.00	0.00	0.00	0.01	0.00	0.01
Num. obs.	12822	12822	12822	12822	12822	12822	12822	12822

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table 4. Each column is a model regressing weather variables on the emotional response. Weather variables were transformed into logistic relationships with smooth thresholds, as described in the text.

This work lays the foundation for a variety of future projects. We now have millions of tweets spanning the US and several months. With these, we plan on exploring the varying effects of weather and extreme events on emotions, using variation over both space and time to identify the effects of weather on emotions while controlling for regional trends in moods. With this and the Sandy data, we also hope to look at the effects of weather on elections, and how those effects are represented in social media.

References

Acar, A. and Muraki, Y. (2011). Twitter for crisis communication: lessons learned from japan's tsunami disaster. *International Journal of Web Based Communities*, 7(3):392–402.

Balog, K., Mishne, G., and de Rijke, M. (2006). Why are they excited?: identifying and explaining spikes in blog mood levels. In *Proceedings of the Eleventh Conference of the European Chapter of the Association for Computational Linguistics: Posters & Demonstrations*, pages 207–210. Association for Computational Linguistics.

Bradley, M. M. and Lang, P. J. (1999). Affective norms for english words (anew): Instruction manual and affective ratings. Technical report, Technical Report C-1, The Center for Research in Psychophysiology, University of Florida.

Heverin, T. and Zach, L. (2010). *Microblogging for Crisis Communication: Examination of Twitter Use in Response to a 2009 Violent Crisis in the Seattle-Tacoma, Washington, Area*. ISCRAM.

Hughes, A. L. and Palen, L. (2009). Twitter adoption and use in mass convergence and emergency events. *International Journal of Emergency Management*, 6(3):248–260.

Jansen, B. J., Zhang, M., Sobel, K., and Chowdury, A. (2009). Twitter power: Tweets as electronic word of mouth. *Journal of the American society for information science and technology*, 60(11):2169–2188.

Kivran-Swaine, F. and Naaman, M. (2011). Network properties and social sharing of emotions in social awareness streams. In *Proceedings of the ACM 2011 conference on Computer supported cooperative work*, pages 379–382. ACM.

Kwak, H., Lee, C., Park, H., and Moon, S. (2010). What is twitter, a social network or a news media? In *Proceedings of the 19th international conference on World wide web*, pages 591–600. ACM.

Mendoza, M., Poblete, B., and Castillo, C. (2010). Twitter under crisis: Can we trust what we rt? In *Proceedings of the first workshop on social media analytics*, pages 71–79. ACM.

Muralidharan, S., Rasmussen, L., Patterson, D., and Shin, J.-H. (2011). Hope for haiti: An analysis of facebook and twitter usage during the earthquake relief efforts. *Public Relations Review*, 37(2):175–177.

Porter, M. F. (1980). An algorithm for suffix stripping. *Program: electronic library and information systems*, 14(3):130–137.

Qu, Y., Wu, P. F., and Wang, X. (2009). Online community response to major disaster: A study of tianya forum in the 2008 sichuan earthquake. In *System Sciences, 2009. HICSS'09. 42nd Hawaii International Conference on*, pages 1–11. IEEE.

Quercia, D., Ellis, J., Capra, L., and Crowcroft, J. (2011). In the mood for being influential on twitter. In *Privacy, security, risk and trust (passat), 2011 ieee third international conference on and 2011 ieee third international conference on social computing (socialcom)*, pages 307–314. IEEE.

Stevenson, R. A., Mikels, J. A., and James, T. W. (2007). Characterization of the affective norms for english words by discrete emotional categories. *Behavior Research Methods*, 39(4):1020–1024.

Tumasjan, A., Sprenger, T. O., Sandner, P. G., and Welpe, I. M. (2010). Predicting elections with twitter: What 140 characters reveal about political sentiment. *ICWSM*, 10:178–185.

Vieweg, S., Hughes, A. L., Starbird, K., and Palen, L. (2010). Microblogging during two natural hazards events: what twitter may contribute to situational awareness. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1079–1088. ACM.