

RESEARCH LETTER

10.1002/2015GL063709

Key Points:

- Aurora can be mapped in real time using Twitter
- Individual tweets can be combined to provide statistically useful data
- Number of tweets can be normalized to provide accurate overall levels

Correspondence to:

N. A. Case,
nathan.a.case@nasa.gov

Citation:

Case, N. A., E. A. MacDonald, M. Heavner, A. H. Tapia, and N. Lalone (2015), Mapping auroral activity with Twitter, *Geophys. Res. Lett.*, 42, 3668–3676, doi:10.1002/2015GL063709.

Received 2 MAR 2015

Accepted 27 APR 2015

Accepted article online 4 MAY 2015

Published online 21 May 2015

Mapping auroral activity with Twitter

N. A. Case^{1,2}, E. A. MacDonald^{1,2}, M. Heavner^{1,3}, A. H. Tapia⁴, and N. Lalone⁴

¹New Mexico Consortium, Los Alamos, New Mexico, USA, ²NASA Goddard Space Flight Center, Greenbelt, Maryland, USA, ³Los Alamos National Lab, Los Alamos, New Mexico, USA, ⁴College of Information Sciences and Technology, Pennsylvania State University, University Park, Pennsylvania, USA

Abstract Twitter is a popular, publicly accessible, social media service that has proven useful in mapping large-scale events in real time. In this study, for the first time, the use of Twitter as a measure of auroral activity is investigated. Peaks in the number of aurora-related tweets are found to frequently coincide with geomagnetic disturbances (detection rate of 91%). Additionally, the number of daily aurora-related tweets is found to strongly correlate with several auroral strength proxies ($r_{\text{avg}} \approx 0.7$). An examination is made of the bias for location and time of day within Twitter data, and a first-order correction of these effects is presented. Overall, the results suggest that Twitter can provide both specific details about an individual aurora and accurate real-time indication of when, and even from where, an aurora is visible.

1. Introduction

With the advent and subsequent rise of social media services, such as Twitter, researchers have been offered an unprecedented level of access to real-time information about events occurring throughout the world, provided by hundreds of millions of users. These users, often termed soft sensors [Tapia *et al.*, 2011], can provide useful information such as their location and the conditions around them.

Studies have shown that Twitter users, who post short updates (140 characters max) known as “tweets,” can provide real-time information about large-scale events and disasters. Examples include earthquakes [Earle *et al.*, 2010; Crooks *et al.*, 2013], influenza outbreaks [Culotta, 2010; Lampos *et al.*, 2010], wildfires [Sutton *et al.*, 2008], and service outages [Motoyama *et al.*, 2010].

As such an event occurs, there is a marked increase in the occurrence of tweets relating to that event. For example, Earle *et al.* [2012] showed that following earthquakes there were often significant localized increases in the number of tweets relating to earthquakes. Indeed, especially in poorly instrumented regions of the world, this Twitter-based detection was generally faster than seismographic-based detection [Earle *et al.*, 2012]. Additionally, since the users’ locations can often be gathered from tweets, either through location-enabled tweets (where the user has opted to share their location) or through location extraction algorithms (see Priedhorsky *et al.* [2014]), researchers are able to visualize the real-time evolution or movement of the disaster [e.g., McDougall, 2011].

This study collates tweets (such as the one shown in Figure 1) and investigates the capability of Twitter for real-time analysis and mapping of an aurora, as has been done with other large-scale events such as natural disasters. As with natural disasters, an aurora is a large-scale natural event which is viewable by, and attracts attention from, the general public. Additionally, although an aurora is frequently visible at high latitudes, a visible aurora at lower latitudes (and thus over larger proportions of the global population) is much less frequent.

A direct comparison of tweet occurrences with visible auroral activity requires widespread, accurate, in situ or ground-based measurements of the aurora. Due to limited data availability, especially at lower latitudes, this is not currently feasible. Instead, the tweet occurrences can be compared to several indices which act as proxies for auroral activity.

2. Data Collection

The tweets used in this study were collected by the Aurorasaurus citizen science project [Tapia *et al.*, 2014] from September 2012 to April 2013. They were delivered by SocialFlow (<http://www.socialflow.com/>) using

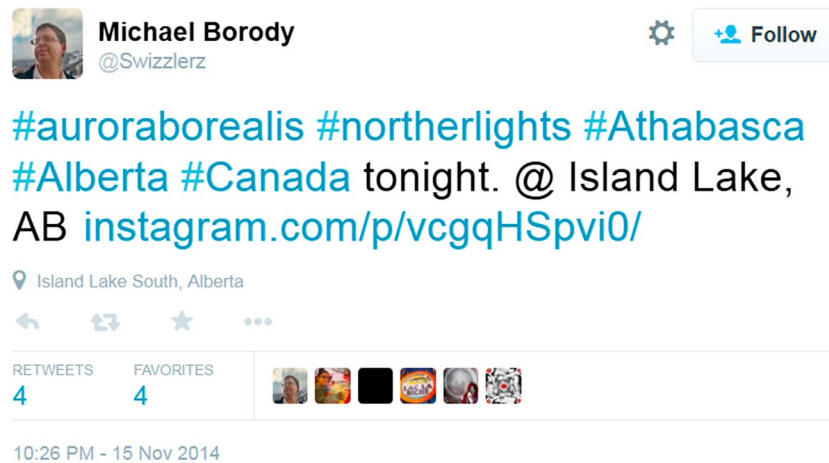


Figure 1. An example of an aurora sighting tweet (<https://twitter.com/Swizzlerz/status/533823437635330050>). Included in this tweet are the time, date, and location of the sighting as well as a link to a photo of the aurora taken by the user.

Twitter’s Firehose streaming service and filtered using aurora-related keywords. The exact filtering criteria varied slightly over time, however, this did not substantially vary the quantity or quality of the returned tweets.

Occasionally, no aurora-related tweets were delivered for a given day. This was the result of a failure in the delivery service rather than there truly being no aurora-related tweets posted. Additionally, several of the days had partial data gaps. The data used in the following study therefore pertains only to those days where the collection service performed as expected throughout the whole day.

In the following analysis, the daily tweet occurrence is compared to several indices which act as auroral strength proxies: Kp [Bartels et al., 1939], Dst [Sugiura, 1964], AE [Davis and Sugiura, 1996], hemispheric power (HP) [Evans, 1987], and the epsilon coupling parameter (ϵ) [Perreault and Akasofu, 1978].

The geomagnetic indices Kp , Dst , and AE all measure disturbances in the Earth’s magnetic field. Kp is a quasi-logarithmic weighted global average of the geomagnetic disturbance at midlatitudes, Dst is an averaged measure of the equatorial magnetic field disturbance (primarily in response to enhancements in the Earth’s ring current), and AE is an average measure of the high-latitude magnetic field disturbance (the result of enhancements in the auroral electrojet).

Hemispheric power (units GW) is an estimate of the total energy deposited into each of the Earth’s hemispheres by precipitating electrons. It is estimated by extrapolating the energy flux measured along the orbital tracks of a constellation of spacecraft. Also calculated in units of power, the epsilon coupling function is an empirically derived parameter that describes the amount of energy being transferred into the Earth’s magnetosphere via an estimate of the electromagnetic Poynting flux. The parameter is sometimes described as a measure of the “solar wind power” [Akasofu, 1981] and, in S.I. units, takes the form

$$\epsilon(W) = \frac{4\pi}{\mu_0} v B_T^2 \sin^4 \left(\frac{\theta}{2} \right) I_0^2 \quad (1)$$

where v is the solar wind velocity, $B_T = \sqrt{B_y^2 + B_z^2}$, θ is the interplanetary magnetic field clock angle ($\theta = \text{atan}(B_y/B_z)$ in GSM co-ordinates) and the empirical scaling factor $I_0 = 9.5R_E$.

The term “proxies” is used since none of the aforementioned are direct measurements of visible auroral activity, rather it is inferred that an increase in the proxies will mean that an aurora is more likely to be visible at lower latitudes and thus viewable by a larger percentage of the global population [e.g., Akasofu and Chapman, 1962; Carbary, 2005].

NASA’s OMNIweb data set (<http://omniweb.gsfc.nasa.gov/>) provides both high- (minute) and low- (hour) resolution solar wind and geomagnetic data. We utilize this data set as a convenient source of the Kp , Dst , and AE indices as well as the solar wind data needed to calculate ϵ . Both Dst and AE are “provisional” and may be subject to small changes once made final. HP is provided by NOAA’s Space Weather Prediction Center using data from the POES satellites.

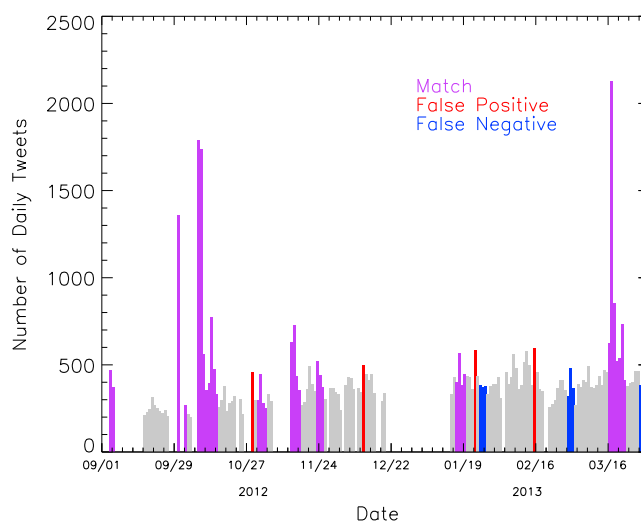


Figure 2. A plot of the number of daily tweets between September 2012 and April 2013. Purple bars indicate that a peak in the daily occurrence of tweets coincided with a geomagnetic disturbance. Red bars indicate that there was a peak in the daily occurrence of tweets but that there was no associated geomagnetic disturbance (i.e., a false positive). Blue bars indicate that a geomagnetic disturbance occurred but there was no corresponding peak in the occurrence of tweets (i.e., a false negative).

The daily tweet occurrence is compared to the daily maximum of each of the proxies. For the purpose of this initial study, the daily values are sufficient (though future studies might look at higher-resolution data). The daily maximum is chosen since we are interested in the peak auroral activity during each day.

3. Results

The Twitter data set used in this study spans from 1 September 2012 to 1 April 2013. Due to the aforementioned data gaps, there are a total of 147 days of data. Locations were associated with $\sim 68\%$ of the tweets, resulting in approximately 62,500 geolocated tweets being collected (an average of ~ 425 per day). For each of the 147 days, the total number of tweets was compared to the daily maximum (minimum for Dst) of each of the aurora proxies (Kp , Dst , AE , HP , and ϵ).

Presented in Figure 2 is a plot of the daily number of aurora-related tweets. Purple bars indicate a peak in the number of tweets which coincides with a period of strong geomagnetic activity (a “match”). If there is a peak in the number of daily tweets but this peak does not coincide with a period of strong geomagnetic activity (a “false positive”) then the bar is red; if there is a strong activity but no peak in the number of tweets (a “false negative”) then the bar is blue.

A peak is defined as when the daily tweet number is at least 30% higher than the 15 day boxcar average background level. Days of strong geomagnetic activity are defined as those days where $Dst < -40$ nT; a somewhat crude filter but sufficient for our purposes [c.f., Russell *et al.*, 1974; Denton *et al.*, 2006]. A “period” of strong geomagnetic activity starts one day preceding the day of strong geomagnetic activity and ends two days following it. Consecutive days of strong geomagnetic activity mean that a period may span for several days.

Of the 11 strong geomagnetic activity periods in this data set, eight were detected through a peak in the number of aurora-related tweets resulting in a detection rate of 73%. Treating each day as individual events, thus allowing us to calculate the number of true negatives, the overall statistical accuracy is found to be 91% with a sensitivity of 79% (see Zhu *et al.* [2010] for definitions).

The mean tweet occurrence during “nonmatch” periods is ~ 250 per day, with a standard deviation of 190. This value can be thought of as an overall “background level”. For “match” periods the mean count is ~ 525 per day (an increase of 110%), with a standard deviation of 500. Both of these distributions have extremely large standard deviations, when compared to the means, which demonstrate that the data are very variable.

In Figure 3 (top row), the number of daily tweets is compared to each of the proxies. The Pearson’s correlation coefficient (r) is shown in each plot and a least squares linear fit is plotted to the data. In Figure 3 (bottom row), the data is filtered to those days when the individual proxies were elevated. As far as we know, there is

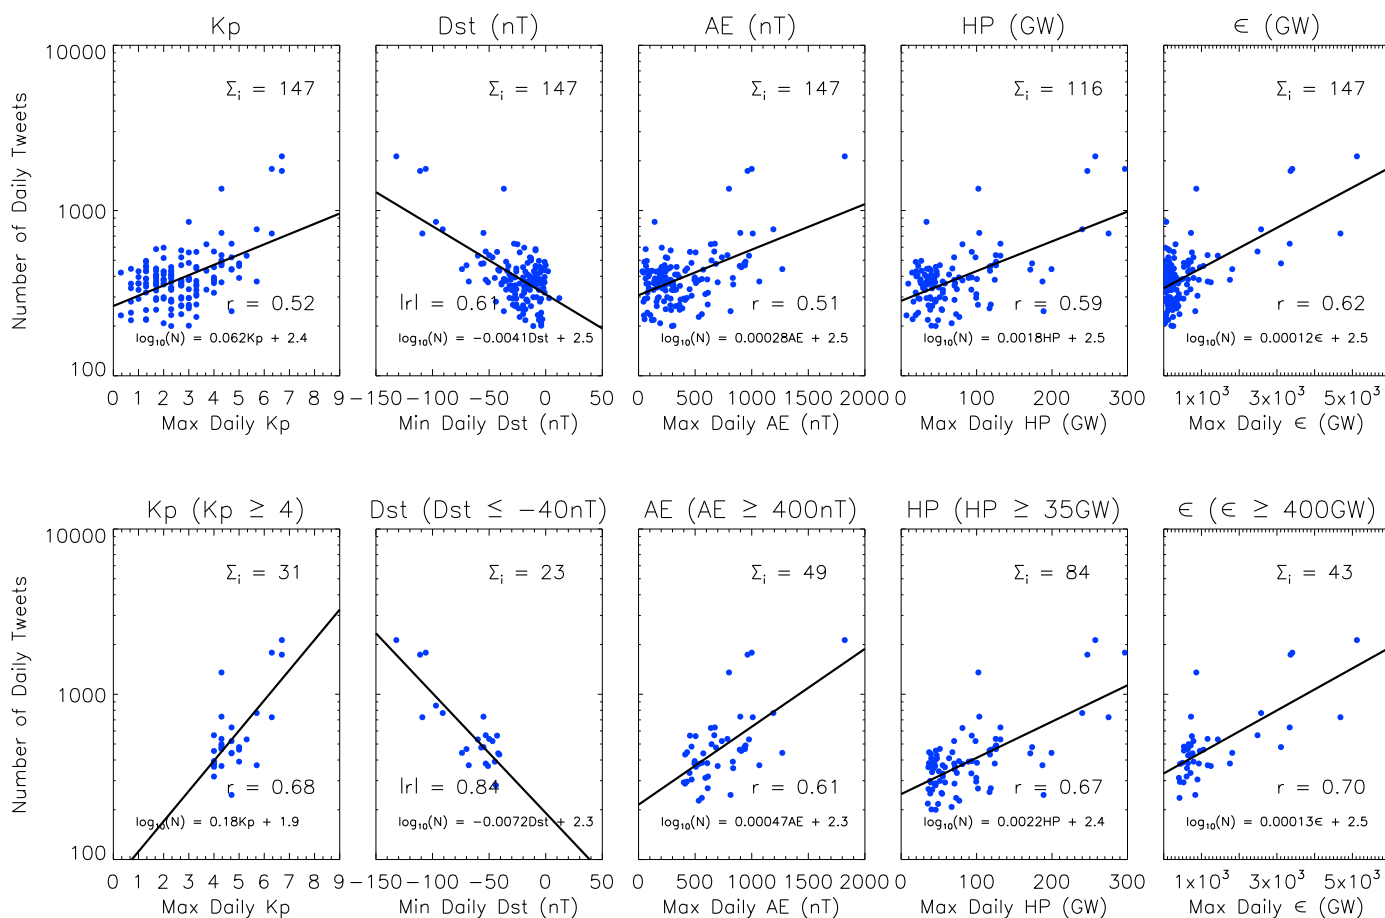


Figure 3. (top row) The daily occurrence of tweets is compared against the auroral proxies (Kp , Dst , AE , HP , and ϵ). (bottom row) The daily occurrence of tweets is compared with the proxies when each was elevated (e.g., $Kp \geq 4$). The Pearson's correlation coefficient is shown in the bottom right of each plot, and the least squares linear line of best fit is plotted.

no strict definition of “elevated” for each of the proxies, however, the following levels seem well used in the literature: $Kp \geq 4$, $Dst \leq -40$ nT, $AE \geq 400$ nT, $HP \geq 35$ GW and $\epsilon \geq 400$ GW. We note that, as shown in Figure 3, these definitions result in different sample sizes for each proxy (especially HP which has approximately twice the amount of data points of the other proxies).

Listed in order of lowest to highest, the correlation coefficients for each of the auroral proxies are AE 0.51, Kp 0.52, HP 0.59, Dst 0.61, and ϵ 0.62. In all cases, correlations are improved when the auroral proxies are elevated: AE 0.61 (a change of +20%), HP 0.67 (+14%), Kp 0.68 (+31%), ϵ 0.70 (+13%), and Dst 0.84 (+38%).

3.1. Data Normalization

Soft sensors from multiple countries might report sightings of the same auroral event. Additionally, since an aurora can occur over a wide range of longitudes and time zones, the same auroral display may only become visible in certain regions hours after being visible in others. Both of these issues complicate using Twitter as an overall measure of auroral activity and suggest that some normalization work should be undertaken.

The number of Twitter users in each country is variable and so even a large rise in aurora-related tweets from a country with a relatively small Twitter user base (e.g., Australia) might not make a significant impact in the overall daily tweet count. In order to account for this, we “normalize” the daily tweets occurrences from the following countries: Australia, Canada, United Kingdom, and United States (these are predominantly English-speaking countries, with at least a moderate chance of viewing the aurora and they have a sizable number of Twitter users in them [Sysomos, 2009]). New Zealand, which has a higher chance of seeing the aurora than Australia, is not included in the normalization list owing to a lack of data on the number of Twitter users living there. The normalization process is to determine the number of tweets originating from each of those countries and multiply that value by an appropriate normalization factor. The first normalization factor

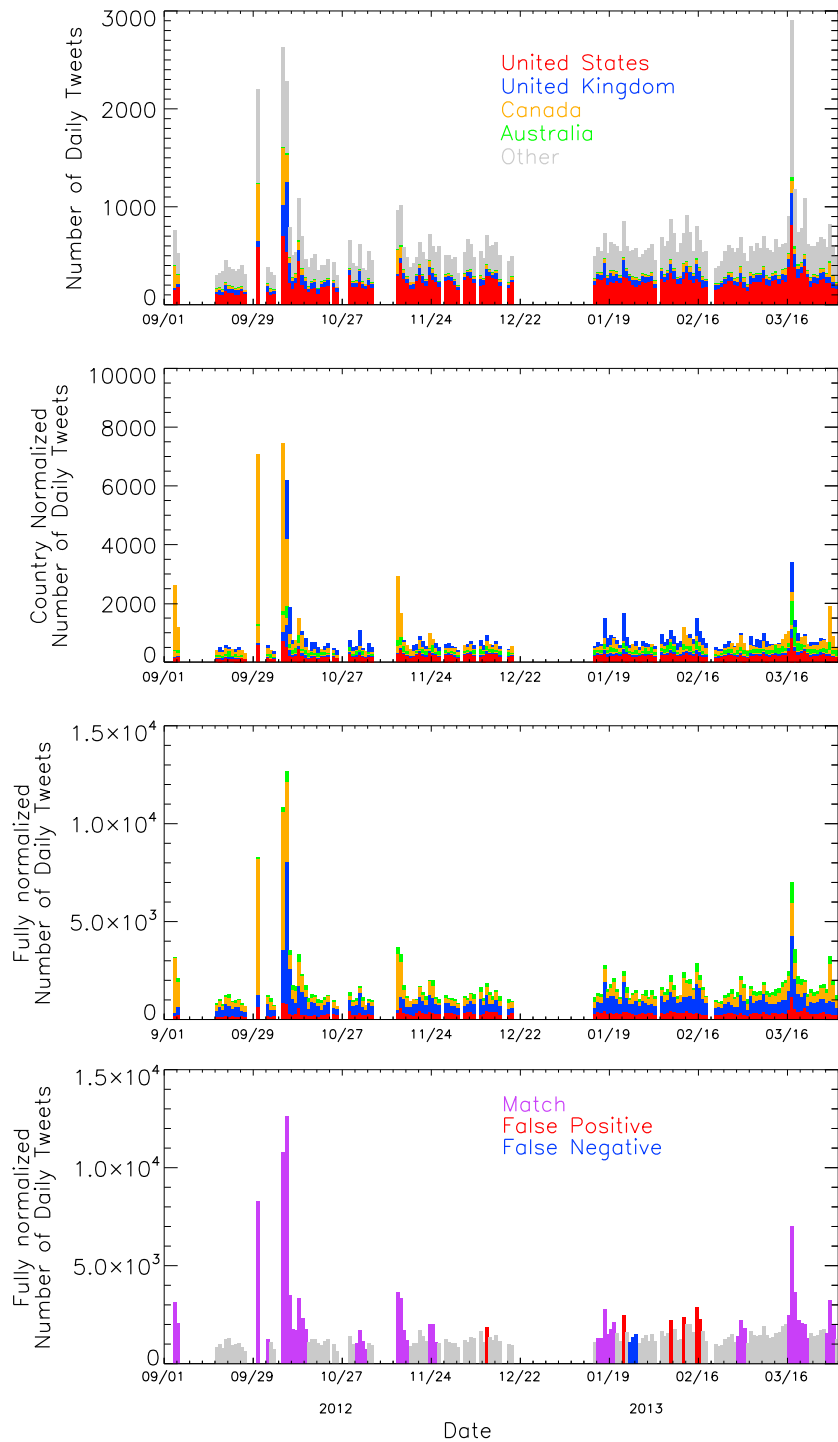


Figure 4. (first panel) The daily occurrences of tweets, colored by country of origin. (second panel) The occurrences are then normalized to account for the number of Twitter users in each of the listed countries. (third panel) The levels are further adjusted to account for hourly variations in the frequency of tweet postings. (fourth panel) These fully normalized daily tweet occurrences are then compared to periods of geomagnetic disturbance (as in Figure 2).

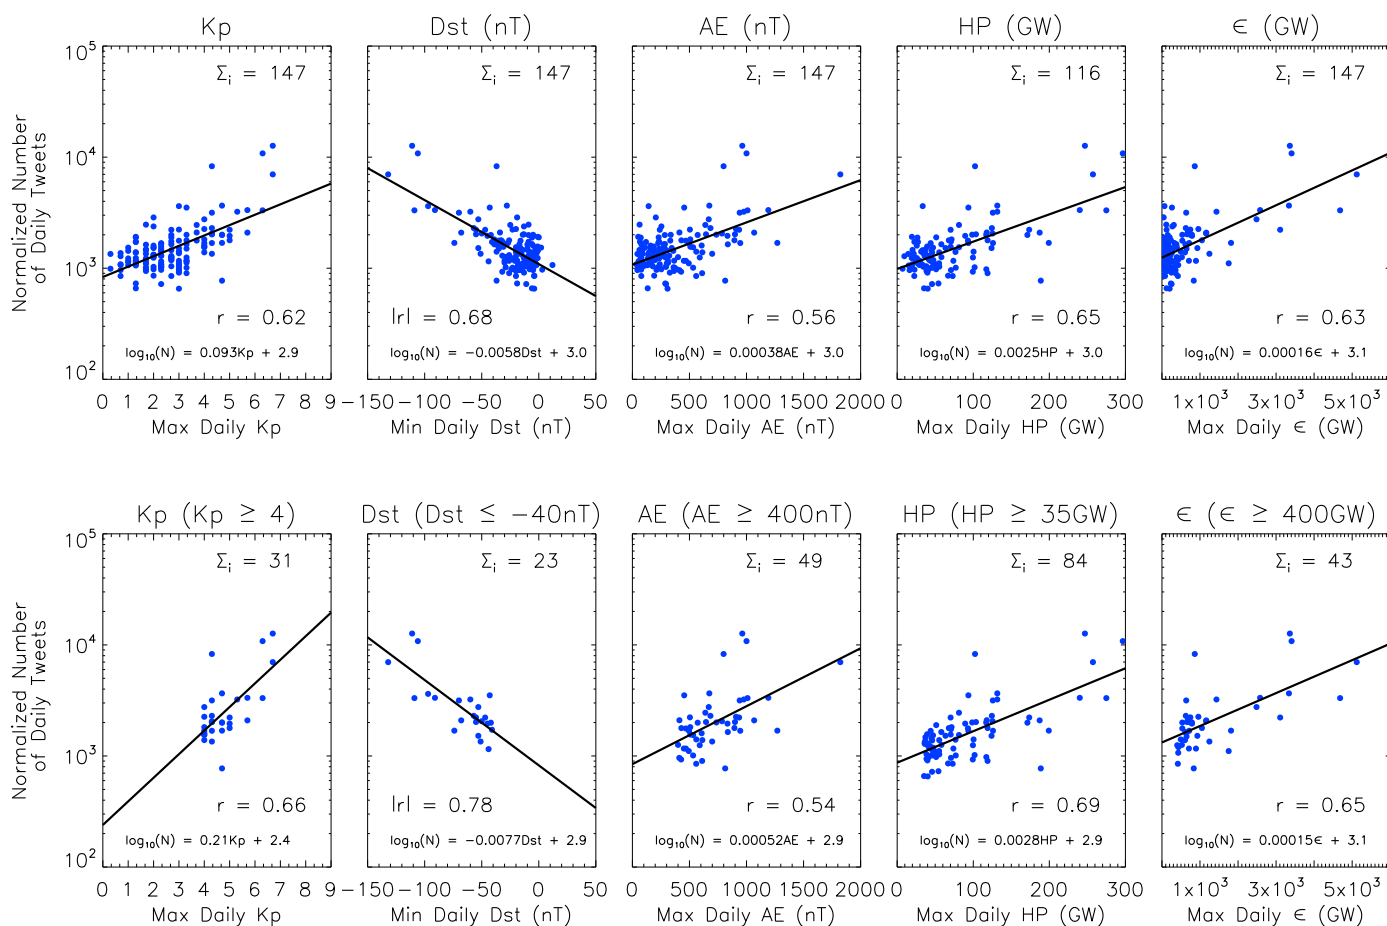


Figure 5. Of the same form as Figure 3, with the normalized daily occurrence of tweets now compared against the auroral proxies.

is determined for each country by calculating the ratio of the number of Twitter users in that country compared to the number in the US.

The frequency of tweets also varies throughout the day. The highest percentage of daily tweets occurs during the afternoon while the lowest percentage occurs during the early hours of the morning [Sysomos, 2009]. Since an aurora is only visible during darkness, adjusting for the local time of day when determining overall tweet levels is therefore also important. The number of tweets per hour is scaled by calculating the ratio of the frequency of tweets in that hour to the maximum frequency.

In Figure 4 (top), the daily number of tweets is plotted. The stacked bars are colored to indicate the country of origin of the tweets. In Figure 4 (second panel), the daily tweet levels are normalized by country. In Figure 4 (third panel), the hourly tweet totals are normalized using the percentages given in Sysomos [2009]. The hourly tweet totals are then summed to produce the “fully normalized” daily tweet occurrences shown. The fully normalized daily occurrence of tweets is compared to periods of strong geomagnetic activity (in the same form as Figure 2) in Figure 4 (fourth panel).

Ten of the 11 periods of strong geomagnetic activity match with peaks in the number of tweets, a detection rate of 91%. Again, treating each day as an individual event, we find that the overall accuracy is now 94% with a sensitivity of 93%.

The normalized daily tweet numbers are then compared with the auroral proxies in Figure 5. The correlation coefficients, listed in order from lowest to highest, are AE 0.56, Kp 0.62, ϵ 0.63, HP 0.65, and Dst 0.68. During elevated periods only, the correlation coefficients are AE 0.54 (−5.3%), ϵ 0.65 (+3.1%), Kp 0.66 (+6.5%), HP 0.69 (+6.2%), and Dst 0.78 (+15%).

4. Discussion

In this study, tweets relating to the aurora, posted between September 2012 and March 2013, were collated and compared to five auroral proxies (Kp , Dst , AE , HP , and ϵ). In Figure 3, the number of daily tweets showed positive correlation with all of the proxies. All correlations were improved when filtering the data to those periods when the proxies were elevated (e.g., $Kp \geq 4$), though both Kp and Dst showed a particularly marked increase ($\sim 35\%$). Using all data, ϵ was found to have the strongest correlation ($r = 0.62$), and using just the data from elevated periods, Dst demonstrated the strongest correlation ($|r| = 0.84$).

The daily tweet occurrence was compared with the daily maximum of each auroral proxy, rather than using hourly values, since the lifetime of an auroral display often spans many hours and the hourly tweet levels were not necessarily consistent throughout the whole period. We expect that this is often due to an “anticipatory” rise in tweets (such as “The aurora will be visible tonight”) or after-the-event tweets (such as “I saw the aurora last night”).

Auroras are seen during the night and so a single event may cause consecutive peaks in the daily number of tweets. Additionally, there will be events that span multiple days (i.e., due to long-spanning storms or reoccurring substorms) and this may help to explain some of the spread in the correlation plots.

To attempt to account for variations in the number of tweets, caused by factors such as the number of Twitter users per country and the time of day, the daily tweet occurrences were normalized. The fairly straightforward attempt at normalization improved the correlation coefficients for all proxies when using all of the tweets (i.e., not filtering to elevated periods). As shown in Figure 5, Dst correlated most strongly with both the unfiltered data ($|r| = 0.68$) and the filtered data ($r = 0.78$). With the exception of AE , filtering the normalized data to periods when the proxies were elevated produced slightly stronger correlations. Interestingly, these correlations were not, however, as strong as with the elevated nonnormalized tweets. Further investigation into why this is the case is warranted and may improve the normalization technique. For example, the normalization values we calculate pertain to data from Sysomos [2009], who studied the whole output of Twitter, and so these values might not be accurate when investigating aurora-related tweets only.

Of course, good correlation coefficients alone do not necessarily show that tweet occurrences can be used to accurately determine when strong auroral displays are occurring. Therefore, in addition to the correlation analysis, the daily tweet levels were compared to periods of strong geomagnetic disturbances (i.e., $Dst < -40$ nT). As shown in Figures 2 and 4, peaks in the number of daily tweets were often found to coincide with periods of strong geomagnetic activity (73% detection rate for nonnormalized and 91% for normalized).

With both the nonnormalized and normalized data, there were false positive events (where a peak in the number of tweets was detected but Dst was not elevated). It is important to note, however, that, although widely used in aurora-related contexts, Dst is actually a measure of the disturbance in the near-equatorial magnetic field due to changes in the Earth’s ring current. As such, Dst acts only as a proxy for auroral strength and not a true measure of it. Auroral displays, especially those at higher latitudes, may occur during smaller Dst values.

Our approach at collecting aurora-related tweets, without manually verifying each of them, is, by definition, imperfect and will result in some level of “spam” tweets. Whether these spam tweets have nothing to do with an aurora (e.g., they mention a town or person called Aurora) or are perhaps not related to current auroral conditions (e.g., “I want to see the aurora”) or are “retweets” of sightings is not accounted for in this study and may explain some of the spread when tweets are compared with the auroral proxies.

Future work, such as that being undertaken by the Aurorasaurus project, should focus on removing these spam tweets by improving filtering techniques. Such improvements may include better keyword or sentiment filtering [e.g., Go *et al.*, 2009]. Additionally, combining automated routines with citizen scientists, who can validate the tweets, could provide more accurate data [e.g., Yu *et al.*, 2012].

5. Conclusions

Twitter has consistently shown to be useful in mapping large-scale events, such as natural disasters, in near real time. An aurora shares several similarities with natural disasters including relative rarity (at least at lower magnetic latitudes), large spatial scales, and strong interest from the general public.

Individual tweets, such as the one shown in Figure 1, can contain information such as the location, a description, or even a photo/video of an aurora. These specific types of tweets can be used to test auroral oval models (i.e., by comparing the location of the sighting and the predicted auroral oval) and provide useful statistics about auroral characteristics (e.g., type, color, and activity level). Combining multiple simultaneous tweets can provide a real-time, large-scale view of where an aurora can currently be seen.

The results of this study show that that peaks in the number of tweets can reliably be used to detect periods of strong geomagnetic activity (as measured by *Dst*). During the time period studied, using the normalized tweet occurrences, 91% of the periods of strong geomagnetic activity were detected (see Figure 4).

Additionally, the results showed that the daily level of tweets increased linearly with increased auroral activity (as inferred from several auroral proxies) and that the two data sets strongly correlate ($r_{\text{avg}} \approx 0.7$). Although the number of tweets was not compared directly with auroral activity, rather just proxies of it, both of these results suggest that Twitter can be used as an accurate measure of the total auroral activity in real time. We note that direct comparisons with auroral activity could perhaps be made, in future case studies, with ground-based all-sky cameras.

Location data (at least to state or country level) were available for a significant percentage of tweets (approximately 68%). Utilizing similar location extraction techniques in real time could therefore suggest from where an aurora is currently visible. This would be useful not only for aurora alert services, who let their users know that an aurora might be visible nearby, but also for comparison with auroral oval models.

Using a fairly straightforward attempt at accounting for geographic and time-of-day biases improved the detection rate of periods of strong geomagnetic activity. Additionally, it also yielded stronger correlations with each of the auroral proxies, at least when all of the data was used (i.e., not filtered to elevated periods). Further enhancements to this normalization technique (for example, investigating the effect of seasonal variance—which seems likely to be important when considering countries from both hemispheres) would seem a worthwhile endeavor.

In summary, during large geomagnetic disturbances, when an aurora is visible at lower latitudes and over highly populated areas, we expect that Twitter will yield a new rich and precise set of data that are highly correlated with and indicative of strong auroral activity. For example, the recent St. Patrick's Day geomagnetic storm, in which the aurora was seen from as far south as central Europe and the mid-U.S., will allow for future analysis of how Twitter responds during an extreme event.

Acknowledgments

This material is based upon work supported, in part, by the National Science Foundation (NSF) under grant 1344296. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of NSF. The OMNI data were obtained from the GSFC/SPDF OMNIWeb interface at <http://omniweb.gsfc.nasa.gov>. The Hemispheric Power data were provided by the National Oceanic and Atmosphere Administration (NOAA) POES satellites and were obtained through the Space Weather Prediction Center (<http://www.swpc.noaa.gov/>).

The Editor thanks an anonymous reviewer for his/her assistance in evaluating this paper.

References

- Akasofu, S.-I. (1981), Energy coupling between the solar wind and the magnetosphere, *Space Sci. Rev.*, 28(2), 121–190.
- Akasofu, S.-I., and S. Chapman (1962), A large change in the distribution of the auroras during the 11 February 1958 magnetic storm, *J. Atmos. Terr. Phys.*, 24(8), 740–742.
- Bartels, J., N. H. Heck, and H. F. Johnston (1939), The three-hour-range index measuring geomagnetic activity, *Terr. Magn. Atmos. Electr.*, 44(4), 411–454, doi:10.1029/TE044i004p00411.
- Carbary, J. F. (2005), A Kp-based model of auroral boundaries, *Space Weather*, 3, S10001, doi:10.1029/2005SW000162.
- Culotta, A. (2010), Detecting influenza outbreaks by analyzing Twitter messages. CoRR, arXiv:1007.4748.
- Crooks, A., A. Croitoru, A. Stefanidis, and J. Radzikowski (2013), # Earthquake: Twitter as a distributed sensor system, *Trans. GIS*, 17(1), 124–147.
- Davis, T. N., and M. Sugiura (1996), Auroral electrojet activity index AE and its universal time variations, *J. Geophys. Res.*, 71, 785–801, doi:10.1029/JZ071i003p00785.
- Denton, M. H., J. E. Borovsky, R. M. Skoug, M. F. Thomsen, B. Lavraud, M. G. Henderson, R. L. McPherron, J. C. Zhang, and M. W. Liemohn (2006), Geomagnetic storms driven by ICME- and CIR-dominated solar wind, *J. Geophys. Res.*, 111, A07507, doi:10.1029/2005JA011436.
- Earle, P., M. Guy, R. Buckmaster, C. Ostrum, S. Horvath, and A. Vaughan (2010), OMG earthquake! Can Twitter improve earthquake response?, *Seismol. Res. Lett.*, 81(2), 246–251.
- Earle, P. S., D. C. Bowden, and M. Guy (2012), Twitter earthquake detection: Earthquake monitoring in a social world, *Ann. Geophys.*, 54(6), 708–715.
- Evans, D. S. (1987), Global Statistical Patterns of Auroral Phenomena, Proceedings of the Symposium on Quantitative Modeling of Magnetospheric - Ionospheric Coupling Processes, 325, Kyoto.
- Go, A., R. Bhayani, and L. Huang (2009), Twitter sentiment classification using distant supervision, CS224N Project Report, Stanford, 1–12.
- Lamos, V., T. De Bie, and N. Cristianini (2010), Flu detector-tracking epidemics on Twitter, in *Machine Learning and Knowledge Discovery in Databases*, pp. 599–602, Springer, Berlin.
- McDougall, K. (2011), Using volunteered information to map the Queensland floods, Proceedings of the 2011 Surveying and Spatial Sciences Conference: Innovation in Action: Working Smarter (SSSC 2011), 13–23, Surveying and Spatial Sciences Institute, Wellington.
- Motoyama, M., B. Meeder, K. Levchenko, G. M. Voelker, and S. Savage (2010), Measuring online service availability using Twitter, WOSN10, 13–13, Boston, Mass.
- Russell, C. T., R. L. McPherron, and R. K. Burton (1974), On the cause of geomagnetic storms, *J. Geophys. Res.*, 79(7), 1105–1109, doi:10.1029/JA079i007p01105.
- Perreault, P., and S.-I. Akasofu (1978), A study of geomagnetic storms, *Geophys. J. R. Astron. Soc.*, 54, 547–573.

- Priedhorsky, R., A. Culotta, and S. Y. Del Valle (2014), Inferring the origin locations of tweets with quantitative confidence, in *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing*, ACM, Baltimore, Md.
- Sugiura, M. (1964), Hourly values of equatorial Dst for the IGY, *Ann. Int. Geophys. Year*, 35, 9.
- Sutton, J., L. Palen, and I. Shklovski (2008), Backchannels on the front lines: Emergent uses of social media in the 2007 southern California wildfires, in *Proceedings of the 5th International ISCRAM Conference*, 624–632, Washington D. C.
- Sysomos (2009), *Inside Twitter: An In-Depth Look Inside the Twitter World (revised 2014)*, 21 pp., Sysomos/Marketwired, Toronto, Canada.
- Tapia, A. H., K. Bajpai, B. J. Jansen, J. Yen, and L. Giles (2011), Seeking the trustworthy tweet: Can microblogged data fit the information needs of disaster response and humanitarian relief organizations, in *Proceedings of the 8th International ISCRAM Conference*, 1–10, Lisbon.
- Tapia, A. H., N. LaLone, E. A. MacDonald, M. Hall, N. A. Case, and M. Heavner (2014), AURORASaurus: Citizen Science, Early Warning Systems and Space Weather, Second AAI Conference on Human Computation and Crowdsourcing, pp. 30–32, doi:10.13140/2.1.4865.0563, Pittsburgh, Va.
- Yu, J., S. Kelling, J. Gerbracht, and W. K. Wong (2012), Automated data verification in a large-scale citizen science project: A case study, 2012 *IEEE 8th International Conference on In E-Science (e-Science)*, pp. 1–8, doi:10.1109/eScience.2012.6404472, Chicago, Ill.
- Zhu, W., N. Zeng, and N. Wang (2010), Sensitivity, specificity, accuracy, associated confidence interval and ROC analysis with practical SAS implementations, NESUG proceedings: Health care and life sciences, 1–9, Baltimore, Md.