A Sensitive Twitter Earthquake Detector

Bella Robinson
CSIRO ICT Centre
G.P.O. Box 664
Canberra, ACT 2601, Australia
bella.robinson@csiro.au

Robert Power
CSIRO ICT Centre
G.P.O. Box 664
Canberra, ACT 2601, Australia
robert.power@csiro.au

Mark Cameron
CSIRO ICT Centre
G.P.O. Box 664
Canberra, ACT 2601, Australia
mark.cameron@csiro.au

ABSTRACT

This paper describes early work at developing an earthquake detector for Australia and New Zealand using Twitter. The system is based on the Emergency Situation Awareness (ESA) platform which provides all-hazard information captured, filtered and analysed from Twitter. The detector sends email notifications of evidence of earthquakes from Tweets to the Joint Australian Tsunami Warning Centre.

The earthquake detector uses the ESA platform to monitor Tweets and checks for specific earthquake related alerts. The Tweets that contribute to an alert are then examined to determine their locations: when the Tweets are identified as being geographically close and the retweet percentage is low an email notification is generated.

The earthquake detector has been in operation since December 2012 with 31 notifications generated where 17 corresponded with real, although minor, earthquake events. The remaining 14 were a result of discussions about earthquakes but not prompted by an event. A simple modification to our algorithm results in 20 notifications identifying the same 17 real events and reducing the *false positives* to 3. Our detector is sensitive in that it can generate alerts from only a few Tweets when they are determined to be geographically close.

Categories and Subject Descriptors

H.4 [**Information Systems Applications**]: Miscellaneous

General Terms

Experimentation, Measurement, Performance

Keywords

Crisis Coordination, Disaster Management, Situation Awareness, Social Media

1. INTRODUCTION

The Joint Australian Tsunami Warning Centre (JATWC) is operated by Geoscience Australia (GA) and the Australian Bureau of Meteorology. GA's role is to detect earthquakes in the oceans around Australia that may be tsunamigenic by using a network of seismic stations. The seismic data is processed automatically to detect earthquakes which are then reviewed by seismologists for final analysis. GA inform the Bureau within 15 minutes of an earthquake occurring. The Bureau manage sea-level monitoring equipment and complex computer models to identify a tsunami when notified of an earthquake event by GA and estimate the impact to the Australian coast. The Bureau also issue tsunami warnings as required.

Identifying earthquakes is a complex task requiring verification from the global network of seismic stations. This process is time consuming due to the need for seismic waves to reach enough seismometers to determine the location and magnitude of earthquakes. For earthquakes that occur close to population centres, first hand reports available on Twitter have been shown to provide a faster means of detection [5, 3]. CSIRO have been working with GA to evaluate the suitability of the ESA platform to identify earthquakes. This paper reports on our early results.

The rest of the paper is organised as follows. First background information is presented about the ESA platform, with a focus on how bursts are detected. Then an overview of the ESA Earthquake Detector is presented followed by an analysis of its operation over a two month period. An outline of planned further work follows and the paper concludes with a summary of our findings.

2. BACKGROUND

2.1 The ESA Platform

The Emergency Situation Awareness (ESA) platform is used to investigate Social Media as a source of information for emergency management. An overview of the technologies developed can be found in [1] while [6] describes the process of identifying Tweets of interest using a burst detector.

Tweets are gathered from the regions shown in Figure 1 using the Twitter search API[1] by providing a search location with geographical coordinates and a search radius. These queries are repeated every twenty seconds to obtain the recent Tweets. As Tweets are obtained from Twitter, they are processed to identify Tweets of interest using a burst detector, described below.

The first tweets collected by ESA were from Queensland during March 2010 with the full captures of Figure 1 starting from September 2011. These captures have been running continuously since then with the occasional interruption due to issues with the IT infrastructure used to gather the Tweets and with problems from Twitter itself. In total,

[1] `https://dev.twitter.com/docs/api/1/get/search`

WWW 2013 Companion, May 13–17, 2013, Rio de Janeiro, Brazil.
ACM 978-1-4503-2038-2/13/05.

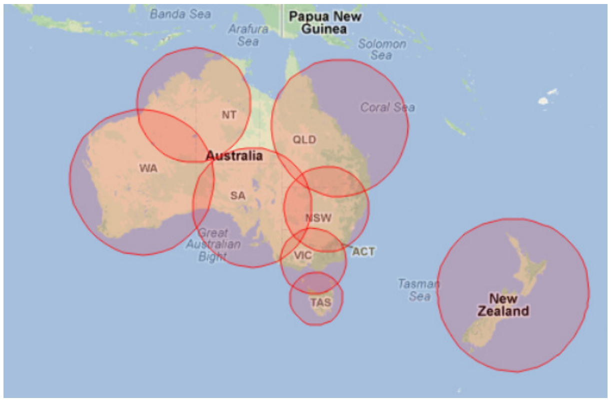

Figure 1: Capture Regions

for the period September 2011–January 2013 over 870 million distinct Tweets have been processed. Duplicate Tweets, approximately 20%, are found due to the overlapping capture regions. A summary of the Tweets processed by region for this period is shown in Table 1.

Table 1: Summary of Tweets Captured By Region

region	Tweets (millions)	Tweets per minute
Qld	159	213
NSW	153	205
Tas	14	19
SA	88	118
WA south	177	237
NT + WA north	175	235
Vic	139	186
NZ	179	240
Total	1084	1453

2.2 ESA Burst Detection

The ESA platform has a burst detector component continuously monitoring variations in feature occurrences coupled with a feature alert map. The ESA burst detection method is based on testing the frequency of observed features in fixed-width time-windows against a statistical content model of historical word frequencies. In essence, the technique identifies situations where an historical model of word frequencies does not fit the observed data.

The burst detector uses a binomial model $B(p_j, N)$ to generate an expected distribution of feature occurrences in a given time window. This scheme only requires a probability p_j for each feature j obtained from historical data to model tweet content. At run-time, for a given feature with n observed occurrences in a window of N Tweets, the cumulative distribution $Prob(obs \geq n | p_j, N)$ gives the p-value of observing n or more Tweets.

The burst detector continuously feeds p-values into the feature alert map which performs the feature-based mapping from a p-value to an alert level (green through red). This configuration enables us to vary the sensitivity of alerts based on statistical cutoffs for different features. Currently we have a simple stratification scheme for each alert level where cutoffs are assigned by order of magnitude from a base cutoff. For this work, we have set the base cutoff for 'earthquak' and '#eqnz' to 0.05 while all other features are set at 0.0001

3. ESA EARTHQUAKE DETECTOR

The detector checks for ESA generated alerts matching earthquake related keywords, tests the currency of the alert, determines if the Tweets producing the alert are close, filters for retweets and generates a notification using an evidence based heuristic. These steps are described below.

3.1 Keywords

The ESA earthquake detector checks for 'earthquak' and '#eqnz' alerts. These are the (stemmed) words most commonly used on Twitter in Australia and New Zealand when reporting about an earthquake. After the large 7.1 magnitude Christchurch earthquake in 2010, the hashtag '#eqnz' became commonly referenced by New Zealanders in earthquake related Tweets. Australians have not adopted an earthquake related hashtag since they experience fewer earthquakes in populated areas.

3.2 Alert Currency

When using ESA to detect an earthquake event, it is important to only consider the first alert generated. As the event unfolds, ESA continues to produce earthquake related alerts, however the nature of the Tweets change over time. The first Tweets describing an earthquake are the ones of interest. An earthquake event is considered finished by the earthquake detector when no ESA earthquake related alerts have been received for 30 minutes.

3.3 GeoSpread

A geo-location process is used to determine if a set of earthquake related Tweets originate from approximately the same location. Around 1% of Tweets collected by ESA are geo-tagged. Instead, the Twitter user's location, defined as part of their 'profile', is used. This location is resolved in 70% of Tweets (retweets excluded) to a geographic location in Australia or New Zealand using Yahoo's GeoPlanet service[2]. This figure improves dramatically to 99% for the Tweets contributing to an ESA earthquake alert. Unfortunately, the assigned location is often imprecise: the location text may be 'Australia' or a state abbreviation.

A simple weighting measure has been used whereby more specific locations, for example a suburb name, are weighted higher than a country or state location. The weights are assigned using the formula below, where the *rank* is defined by the Yahoo GeoPlanet service[2]:

$$weight = 15 - rank$$

For example, the location 'Australia' has a *rank* of 14 and is assigned a *weight* of 1. 'New Zealand', a rank of 11, has a *weight* of 4, the city of 'Christchurch' (7) has a *weight* of 8 and the suburb 'Sumner' (2) has a *weight* of 13. Geotagged Tweets are assigned the maximum weighting of 15. These weightings are combined with the Tweet counts for each distinct location to produce a *heatmap*. This provides a visual indication of where the earthquake reports have come from. Note that retweets are excluded from this process.

In addition to producing a *heatmap*, a measure of the geographic spread of the Tweets is generated. This helps determine if an ESA earthquake alert contains first hand reports of an actual event. The following *GeoSpread* measure has been defined for this:

[2]http://developer.yahoo.com/geo/geoplanet/

$$GeoSpread = \frac{\sigma_{WeightedLatitudes} + \sigma_{WeightedLongitudes}}{2}$$

To construct the weighted latitude and longitude value sets, for each distinct location the latitude and longitude are added to the appropriate set (number of Tweets with that location $\times$ location weight) times.

3.4 Alert Heuristic

The *GeoSpread* measure was developed to help characterise the location of a collection of Tweets. The ESA burst detector is based on the statistical analysis of the frequency of (stemmed) words within Tweets. An 'earthquake' burst does not necessarily mean that it relates to an earthquake event. An analysis of previous ESA earthquake alerts revealed that the *GeoSpread* measure and the retweet percentage can be used as effective indicators. A retweet, a re-posting of someone else's Tweet, is not useful for the aim of identifying first hand reports.

To determine the heuristics to use, 12 months of historical ESA earthquake alerts were analysed and a geospread measure of less than 4 and a retweet percentage of less than 18 were determined as suitable. This *GeoSpread* value indicates that Tweets are from the same region. Higher values imply the information is second hand, for example comments on news reports of an earthquake not local to them. A low retweet percentage is required since we are only interested in first hand reports.

3.5 Notification

When an ESA earthquake burst has been identified as described above, an email is sent to the JATWC. An example is shown in Figure 2. The email summarises the information that triggered the notification and includes a hyperlink to the ESA Alert Monitor web interface, Figure 3. This interface allows the user to view the *heatmap* and monitor ongoing developments related to the earthquake, such as reports of damage. Note that expletives have been blurred in these images.

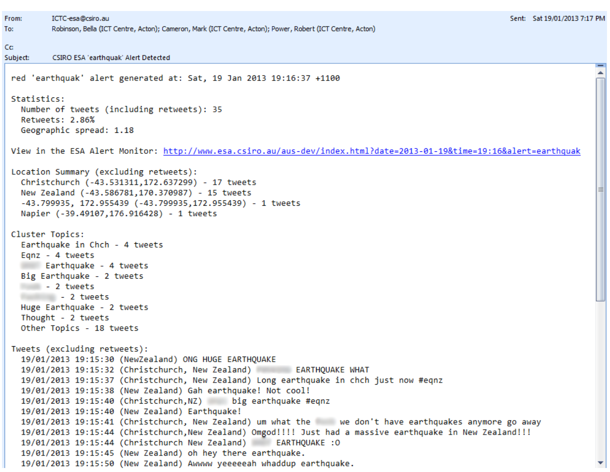

From: ICTC-esa@csiro.au Sent: Sat 19/01/2013 7:17 PM
To: Robinson, Bella (ICT Centre, Acton); Cameron, Mark (ICT Centre, Acton); Power, Robert (ICT Centre, Acton)
Cc:
Subject: CSIRO ESA 'earthquak' Alert Detected

```
red 'earthquak' alert generated at: Sat, 19 Jan 2013 19:16:37 +1100

Statistics:
  Number of tweets (including retweets): 35
  Retweets: 2.86%
  Geographic spread: 1.18

View in the ESA Alert Monitor: http://www.esa.csiro.au/aus-dev/index.html?date=2013-01-19&time=19:16&alert=earthquak

Location Summary (excluding retweets):
  Christchurch (-43.531311,172.637299) - 17 tweets
  New Zealand (-43.586781,170.370987) - 15 tweets
  -43.799935, 172.955439 (-43.799935,172.955439) - 1 tweets
  Napier (-39.49107,176.916428) - 1 tweets

Cluster Topics:
  Earthquake in Chch - 4 tweets
  Eqnz - 4 tweets
  [illegible] Earthquake - 4 tweets
  Big Earthquake - 2 tweets
  [illegible] - 2 tweets
  [illegible] - 2 tweets
  Huge Earthquake - 2 tweets
  Thought - 2 tweets
  Other Topics - 18 tweets

Tweets (excluding retweets):
  19/01/2013 19:15:30 (NewZealand) ONG HUGE EARTHQUAKE
  19/01/2013 19:15:32 (Christchurch, New Zealand) [illegible] EARTHQUAKE WHAT
  19/01/2013 19:15:37 (Christchurch, New Zealand) Long earthquake in chch just now #eqnz
  19/01/2013 19:15:38 (New Zealand) Gah earthquake! Not cool!
  19/01/2013 19:15:40 (Christchurch,NZ) [illegible] big earthquake #eqnz
  19/01/2013 19:15:40 (New Zealand) Earthquake!
  19/01/2013 19:15:41 (Christchurch, New Zealand) um what the [illegible] we don't have earthquakes anymore go away
  19/01/2013 19:15:44 (Christchurch,New Zealand) Omgod!!!! Just had a massive earthquake in New Zealand!!!
  19/01/2013 19:15:44 (Christchurch New Zealand) [illegible] EARTHQUAKE :O
  19/01/2013 19:15:45 (New Zealand) oh hey there earthquake.
  19/01/2013 19:15:50 (New Zealand) Awwww yeeeeeah whaddup earthquake.
```

Figure 2: Example Email

All ESA alerts are shown on the Alert Monitor interface as a coloured tag cloud at the top of the screen. In the example of Figure 3, the red 'earthquak' alert has been selected. The Tweets contributing to the alert are then clustered to identify topics, shown below the alerts, as described in [6]. The Tweets themselves are also displayed so that their contents can be reviewed. The *heatmap* for the selected alert is shown on the right with the calculated *GeoSpread* measure and retweet percentages displayed directly above it.

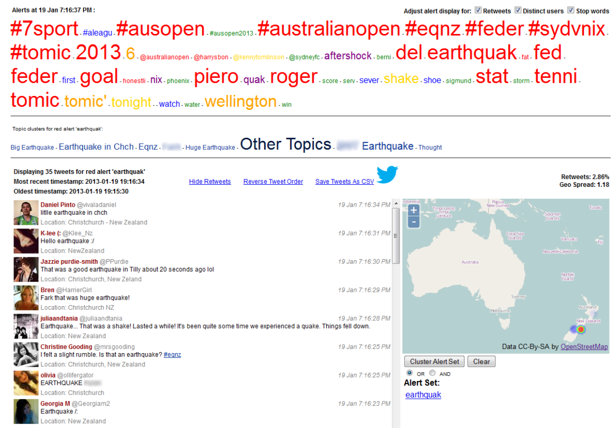

Figure 3: Example Alerts

4. NOTIFICATION ANALYSIS

4.1 Initial Results

The ESA earthquake detector has been operating since 12 December 2012. To analyse its performance, each earthquake related burst was reviewed to determine if an email was sent and the contributing Tweets examined to determine if the burst was related to first hand reports of an earthquake. The results are shown in Table 2.

Table 2: Raw Results.

	email	no email
reports	17 (*TP*)	4 (*FN*)
no reports	14 (*FP*)	64 (*TN*)

T = *True*, *F* = *False*, *P* = *Positive*, *N* = *Negative*

A measure of the accuracy[3] of these results is the *precision* of 0.55, *recall* of 0.81 and an *F1 score* of 0.65, where:

$$precision = \frac{TP}{TP + FP}$$

$$recall = \frac{TP}{TP + FN}$$

$$F1 = 2 \times \frac{precision \times recall}{precision + recall}$$

4.2 Revised Results

These results indicate that the earthquake detector is *too* sensitive. The instances where an email was incorrectly sent were examined to determine if there were features in common. In all of these cases the number of contributing Tweets was low. An additional constraint was imposed: the number of contributing Tweets must be greater than three before an email notification is generated. The revised results are shown in Table 3 and the accuracy measures are now: a *precision* of 0.85, *recall* of 0.77 and an *F1 score* of 0.81.

Note that the total number of alerts in Tables 2 and 3 differ by 1. This is due to one case where the modified

[3] http://en.wikipedia.org/wiki/F1_Score

detector found a burst for 'earthquak' that did not result in a notification (*FN*) due to too few Tweets. However, a subsequent burst for '#eqnz' passed the new condition (*TP*). Whereas the original burst detector recorded only one event, the modified one found two.

Table 3: Revised Results.

	email	**no email**
reports	17 (*TP*)	5 (*FN*)
no reports	3 (*FP*)	75 (*TN*)

4.3 Verification

All of the true positive cases were examined and found to correlate to seismically verified earthquakes as listed by New Zealand's GeoNet[4] and GA[5]. For the true positive cases the average time delay between the earthquake origin time as computed by GeoNet or GA and when ESA sent a notification email was 3:03 (minutes:seconds), with a minimum delay of 1:05 and a maximum delay of 5:34. This compares favourably with NEIC-like organisations that have average seismic-location solution times of between 8 and 12 minutes[6].

It should be noted that for two of our false negative cases, while the Tweets appeared to contain first hand earthquake reports, no corresponding earthquake could be found within a suitable time frame. One is a report of being woken up by an earthquake 'last night', however when looking at the corresponding seismic data, this occurred over 24 hours before the Tweet was created.

During the evaluation period, no severe earthquakes occurred within Australia or New Zealand. The earthquakes reported were between 2.2 and 5.2 in magnitude.

5. FURTHER WORK

Further work is in progress to monitor Tweets beyond Australia and New Zealand, targeting the plate boundary of tsunami source zones as shown in Figure 4. Configuring a burst detector is straightforward once representative data for this region is obtained, though deeper understanding of the many languages used in this region will be challenging. Initial monitoring indicates that the volume of Tweets, currently over 8K per minute, is dramatically larger than what is currently being processed by ESA and may present new challenges for our infrastructure and architecture.

There are other areas to explore also. Determining the location of the Tweet's creator is fundamental to our detector and a combination of methods are being explored to increase the precision of the *GeoSpread* measure. The user's declared profile location is being compared with official place names defined in national Gazetteers and the Tweet content is being used to determine if region specific language cues are available to help disambiguate the user's location [2, 4].

Also, preliminary classifiers have been developed [6], trained on Tweets obtained from the various Christchurch earthquakes. The aim is to follow up the automatic alert email notification with an impact analysis email shortly afterwards, for example 15 minutes later. The impact analysis will provide evidence of impact as reported in Tweets as identified by the classifiers.

[4] `http://www.geonet.org.nz/quakes/`

[5] `http://www.ga.gov.au/earthquakes/home.do`

[6] Daniel Jaksa, personal communication, 21 March 2013.

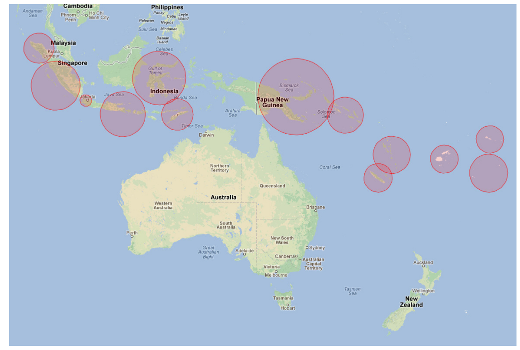

Figure 4: Regional Captures

6. CONCLUSIONS

The ESA platform provides all-hazard information captured, filtered and analysed from Twitter using a background language model to characterise the 'normal' activity. Unusual events can be identified when the observed activity varies from that historically recorded. This can be used as an alert monitor for words of interest.

The ESA platform has been successfully deployed to target Tweets describing earthquake events. Our initial exploratory results were successful in identifying all Tweets mentioning earthquakes however upon review the detector was *too* sensitive. A simple modification to the algorithm to test the number of Tweets contributing the alert increased the accuracy, in terms of the *F1 score*, from 0.65 to 0.81.

Acknowledgments

Thanks go to Daniel Jaksa and the JATWC GA team for their help in developing the ESA earthquake detector.

7. REFERENCES

[1] M. A. Cameron, R. Power, B. Robinson, and J. Yin. Emergency situation awareness from twitter for crisis management. In *Proceedings of the 21st international conference companion on World Wide Web*, WWW '12 Companion, pages 695–698, New York, NY, USA, 2012. ACM.

[2] Z. Cheng, J. Caverlee, and K. Lee. You are where you tweet: a content-based approach to geo-locating twitter users. In J. Huang, N. Koudas, G. J. F. Jones, X. Wu, K. Collins-Thompson, and A. An, editors, *CIKM*, pages 759–768. ACM, 2010.

[3] P. S. Earle, D. C. Bowden, and M. Guy. Twitter earthquake detection: earthquake monitoring in a social world. *Annals of GeoPhysics*, 54(6):708–715, 2012.

[4] Y. Ikawa, M. Enoki, and M. Tatsubori. Location inference using microblog messages. In *Proceedings of the 21st international conference companion on World Wide Web*, WWW '12 Companion, pages 687–690, New York, NY, USA, 2012. ACM.

[5] T. Sakaki, M. Okazaki, and Y. Matsuo. Earthquake shakes twitter users: real-time event detection by social sensors. In *Proceedings of the 19th international conference on World wide web*, WWW '10, pages 851–860, New York, NY, USA, 2010. ACM.

[6] J. Yin, A. Lampert, M. Cameron, B. Robinson, and R. Power. Using social media to enhance emergency situation awareness. *IEEE Intelligent Systems*, 27(6):52–59, 2012.