



**QUEEN'S
UNIVERSITY
BELFAST**

A Machine Learning Analysis of Twitter Sentiment to the Sandy Hook Shootings

Wang, N., Varghese, B., & Donnelly, P. (2017). A Machine Learning Analysis of Twitter Sentiment to the Sandy Hook Shootings. In *2016 IEEE 12th International Conference on e-Science (e-Science)* Institute of Electrical and Electronics Engineers Inc.. <https://doi.org/10.1109/eScience.2016.7870913>

Published in:

2016 IEEE 12th International Conference on e-Science (e-Science)

Document Version:

Peer reviewed version

Queen's University Belfast - Research Portal:

[Link to publication record in Queen's University Belfast Research Portal](#)

Publisher rights

© 2016 IEEE.

This work is made available online in accordance with the publisher's policies. Please refer to any applicable terms of use of the publisher.

General rights

Copyright for the publications made accessible via the Queen's University Belfast Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The Research Portal is Queen's institutional repository that provides access to Queen's research output. Every effort has been made to ensure that content in the Research Portal does not infringe any person's rights, or applicable UK laws. If you discover content in the Research Portal that you believe breaches copyright or violates any law, please contact openaccess@qub.ac.uk.

Open Access

This research has been made openly available by Queen's academics and its Open Research team. We would love to hear how access to this research benefits you. – Share your feedback with us: <http://go.qub.ac.uk/oa-feedback>

A Machine Learning Analysis of Twitter Sentiment to the Sandy Hook Shootings

Nan Wang, Blesson Varghese
School of EECS

Queen's University Belfast, UK
Email: {nwang03, varghese}@qub.ac.uk

Peter D. Donnelly
Dalla Lana School of Public Health
University of Toronto, Canada
Email: peter.donnelly@oahpp.ca

Abstract—Gun related violence is a complex issue and accounts for a large proportion of violent incidents. In the research reported in this paper, we set out to investigate the pro-gun and anti-gun sentiments expressed on a social media platform, namely Twitter, in response to the 2012 Sandy Hook Elementary School shooting in Connecticut, USA. Machine learning techniques are applied to classify a data corpus of over 700,000 tweets. The sentiments are captured using a public sentiment score that considers the volume of tweets as well as population. A web-based interactive tool is developed to visualise the sentiments and is available at <http://www.gunsontwitter.com>. The key findings from this research are: (i) There are elevated rates of both pro-gun and anti-gun sentiments on the day of the shooting. Surprisingly, the pro-gun sentiment remains high for a number of days following the event but the anti-gun sentiment quickly falls to pre-event levels. (ii) There is a different public response from each state, with the highest pro-gun sentiment not coming from those with highest gun ownership levels but rather from California, Texas and New York.

I. INTRODUCTION

On the morning of 14th December 2012 a heavily armed twenty-year-old man with “*significant mental health issues*” [1] and access to his mother’s legally owned firearms shot his way into the locked building of the Sandy Hook Elementary School (SHES) in Connecticut, USA and in less than eleven minutes murdered 20 children and 6 adults. Before driving to the school he had shot and killed his mother in her bed and at the conclusion of the shootings he took his own life.

The world media response to this tragedy was unprecedented in its scale and much has subsequently been written about the events themselves and how future similar events could be avoided [2]. Predictably much attention focussed on issues of gun ownership and control with those advocating tighter gun control clashing with those intent on exercising and protecting (as they saw it) the US constitutional second amendment right to bear arms. That debate continues and, in its totality, is beyond the scope of this paper in which we largely restrict ourselves to seeking to understand and interpret pro-gun and anti-gun sentiment as expressed on a social media platform, namely Twitter in a forty day period starting seven days before the shootings.

In order to achieve this we have used the application of machine learning to a sample of over 700,000 tweets made by individuals in the United States that contain one or more predetermined relevant key words. What we sought

to capture was pro-gun and anti-gun sentiment and how it changed over time. For this we developed a framework to collect, pre-process and classify tweets and further visualise the sentiments. A hand curated gold standard dataset is employed in classifying the entire sample. A number of machine learning approaches are evaluated and the approaches that provide maximum accuracy over a small sample are used for classifying the entire tweet sample. We visualise this sentiment by state taking into account state population size and gun ownership levels through an application we have hosted at <http://www.gunsontwitter.com>.

The key observations are that there is a peak of both anti-gun and pro-gun sentiment on the day of the shooting. Anti-gun sentiment quickly falls to pre-event levels whereas pro-gun sentiment remains high for a number of days. It is interesting to note that all states do not respond in the same way; the highest gun ownership states are not the most pro-gun, but California, Texas and Florida have highest pro-gun sentiment even when their larger population size is taken into account.

The remainder of this paper is organised as follows. Section II presents the methodology adopted in this paper to collect, pre-process, classify, summarise and visualise the sentiment of the tweets. Section III considers the machine learning approaches employed for classifying the tweets. Section IV takes the Sandy Hook school shooting case study into account by applying the methodology on a volume of tweets and visualising pro-gun and anti-gun sentiments. Section V presents the public health implications of the research reported in this paper. Section VII concludes this paper.

II. METHODOLOGY

The methodology for analysing public sentiment incorporates machine learning and (1) collects, (2) pre-processes, (3) classifies, (4) summarises and (5) visualises data. Figure 1 illustrates the sequence of steps in the methodology.

Raw data from Twitter is collected and pre-processed into a trimmed data set in an appropriate format, which is further split into training and testing data sets. In the machine learning stage different algorithms, such as Support Vector Machines (SVM), Maximum Entropy (ME), Tree, Bagged Tree, Boosted Tree, Random Forest (RF), Neural Network (NN) and Naïve Bayes (NB), are explored to build classifiers. A subset of more accurate classifiers are used for classifying the tweets.

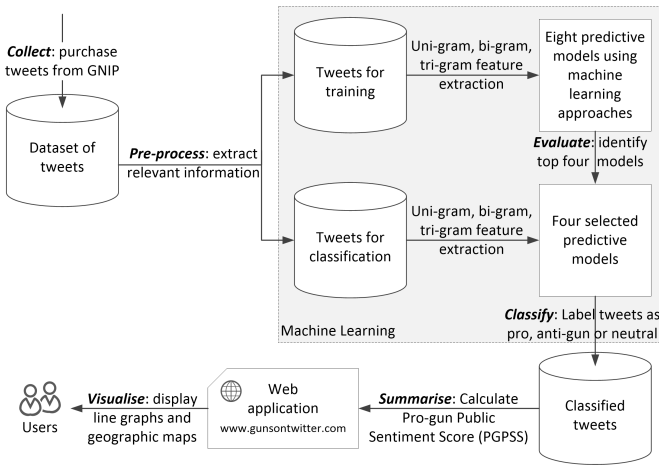


Fig. 1. Methodology for analysing and visualising public sentiment

A number of functions that can be used to derive meaningful information are applied on the classified tweets. The results are summarised using multiple visualisation techniques.

Step 1 - Collect: The raw data of tweets in JSON format is purchased from GNIP¹, a third-party authorised to resell historical tweets determined by predefined rules that takes into account filters based on the time-period, key words and geographic location. Compared to alternatives such as custom-built web crawlers or programming the Twitter Streaming API², this approach saved time and provided data in a consistent format, although the tweets had to be purchased.

Step 2 - Pre-process: A file containing a list of JSON files and their remote locations was obtained from the reseller. Each JSON file is organised for a ten minute period and contains the tweets and a large amount of information regarding the tweet³. The JSON files were parsed to extract the most relevant data for the research and were converted into the CSV format. All further tasks were performed using the R programming language in RStudio⁴. This was preferred to exploit the readily available packages for data mining and natural language processing.

Step 3 - Classify: In the machine learning stage, the sentiments of the tweets are extracted. The training data set contains 5000 tweets selected from the trimmed data while the remaining tweets are left for classification. The sentiment of each tweet is manually labelled into three classes, namely; pro-gun, anti-gun and neutral, to obtain a gold standard dataset. It is assumed that the distribution of tweets is not equal among these classes on grounds of the characteristic of Twitter - a convenient platform for users to express subjective ideas (which convey a pro-gun or anti-gun sentiment). Therefore, the numbers of tweets for three classes are adjusted as 2000 respectively for the pro-gun and anti-gun classes and 1000 for the neutral class. Intentionally giving extra weight to subjective

classes is known to produce classifiers that are more sensitive to these classes [3].

Using the extracted features classifiers are built using supervised training approaches which are presented in the next section. Prior to the deployment of more accurate classifiers, the effectiveness of all models are evaluated by the output of sample accuracy in a 10-fold cross validation [4]. Only top classifiers are selected and further developed based on the features extracted for the classification of the remaining tweets.

Step 4 - Summarise: In addition to the time stamps and coordinates that the data originally contains, a series of methods for calculating a Pro-Gun Public Sentiment Scores (PGPSS) are utilised for comparison. Consider a geographic region defined by g and a time frame t . The baseline PGPSS is denoted as

$$PGPSS_1 = \frac{\text{count}_{(g,t)}(\text{positive tweets})}{\text{count}_{(g,t)}(\text{negative tweets})}$$

Baseline PGPSS measures the number of pro-gun tweets over anti-gun tweets from a set of tweets originating from a state in a given time frame. It gives the degree of positivity in the selected tweets as an index which measures the positive tweets over negative tweets. However, $PGPSS_1$ is not the best indicator of positivity since it does not take the volume of tweets that originates from g . This poses a problem which is illustrated in the following example.

Consider the country g has population of 11 million, which has two states g_1 and g_2 . Suppose the population of g_1 is 10 million and g_2 is 1 million, and the total number of available tweets from the states be 1 million (200,000 positive and 800,000 negative) and 10,000 (8,000 positive and 2,000 negative) respectively. Now $PGPSS_1$ for g_1 is 0.25 and for g_2 is 4 which indicates a higher positive sentiment for g_2 than g_1 . The normalised $PGPSS_1$ for g_1 is 0.0625 and for g_2 is 1. However, this is not entirely accurate in that the volume of tweets that originated from g_2 is relatively a lot smaller than the tweets from g_1 . Hence, a score that takes the volume of tweets that originates from g_1 and g_2 into account is useful.

Hence, we define a PGPSS score, which is corrected for volume of tweets defined as

$$PGPSS_2 = \frac{\text{count}_{(g,t)}(\text{positive tweets})}{\text{count}_{(g,t)}(\text{negative tweets})} * \frac{\text{count}_{(g,t)}(\text{tweets})}{\text{count}_{(t)}(\text{tweets})}$$

The volume correction ratio for g_1 is $1,000,000/1,010,000 = 0.990099$ and for g_2 is 0.009901 . Now $PGPSS_2$ for g_1 is 0.247525 and for g_2 is 0.039604, and the normalised $PGPSS_2$ for g_1 is 1 and for g_2 is 0.16.

An additional correction for population is considered by defining a different PGPSS score which is denoted as

$$PGPSS_3 = \frac{\text{count}_{(g,t)}(\text{positive tweets})}{\text{count}_{(g,t)}(\text{negative tweets})} * \frac{\text{count}_{(g,t)}(\text{tweets})}{\text{count}_{(t)}(\text{tweets})} * \frac{\text{population}_g}{\text{population}}$$

¹<http://gnip.com/>

²<https://dev.twitter.com/docs/streaming-apis>

³<http://gnip.com/sources/twitter/historical/>

⁴<http://www.rstudio.com>

The population correction ratio for g_1 is $1\text{million}/11\text{million} = 0.90909$ and for g_2 is $1000/11\text{million} = 0.000909$. The $PGPSS_3$ score for g_1 is 0.225023 and for g_2 is 0.000036, and the normalised $PGPSS_3$ for g_1 is 1 and for g_2 is 0.000016.

All PGPSS scores are normalised between 0 and 1 for the purpose of visualisation. To interpret the PGPSS scores in the case of public sentiment towards gun use, g_2 is more supportive of guns than g_1 according to $PGPSS_1$, but after correcting the scores, as should be g_1 is presented as a state more supportive of guns than g_2 . The corrections applied to the public sentiment scores provide a better overview in the context of the volume of tweets and population.

Step 5 - Visualise: A web application is developed under the Shiny framework⁵ for visualising data. Motion charts, line graphs and geographic maps are generated using Google Charts API⁶ through the googleVis package [5]. These techniques present results at the country and state levels as well as hourly and daily analyses.

Motion charts display a continuous bubble view of data for a combination of six variables simultaneously; for example, the x and y axes, a bubble, its size and colour and a slider all display relevant information. This technique provides an animated overview of the changes through an interface to observe variations over time. Additionally, a bar plot and a line graph are embedded in the motion chart.

The line graph presents the volume of classified tweets over the dimension of time. A control gadget is available for users to monitor the data on daily or hourly basis. Options on zooming and changing time period are embedded in the graph.

The geographic map visualises a series of PGPSS results in the form of a state-level map. A particular type of Public Sentiment Score and time frame are selected by users each time to be projected into a gradient colour on the map. This technique is useful to interpret the statistics from a geographic perspective.

III. MACHINE LEARNING APPROACHES

Feature extraction and modelling are crucial steps in any machine learning approach. Various techniques on feature selection have been examined in previous research, from the baseline approach of N-gram features [6] to the advanced linguistic analysis on POS tags [7]. Meanwhile, SVM and NB are the top two frequently used algorithms for sentiment analysis and favoured by many researchers due to their superior performance [8]. The evaluation of model performances compares their accuracy from the aspect of training data sizes, N-gram features and algorithms.

A. Feature Extraction

In the feature extraction stage, uni-gram, bi-gram and tri-gram are considered in the domain of N-gram features [6]. As the most frequently used method in text classification, feature extraction using uni-grams process each word within

the sentence individually. Phrases of two and three words are extracted as Bi-gram and tri-gram features for contrasts.

Hashtags and reply/mention tags within each tweet are used as additional features in that sometimes they contain strong subjective opinions. For example in the case of sentiment towards gun use, *#2ndamendment* may indicate pro-gun views on the use of guns whereas *#guncontrol* may suggest more neutral or anti-gun views. Similarly, *@NRA* may be a strong indicator of emotional text, albeit vague in terms of polarity.

B. Modelling

Eight predictive models are developed using the training data set obtained from machine learning approaches, namely Support Vector Machine (SVM), Naïve Bayes (NB), Maximum Entropy (ME), Tree, Bagging, Boosting, Random Forest (RF) and Neural Network (NN). With the exception of Nave Bayes (NB), which is developed with e1071 package [9], all the other classifiers are built with RTextTools [10], a machine learning library for text classification.

SVM uses kernels to find a hyperplane that divides data into different categories with the maximum margin [11, chapter 9]. In RStudio this model is built with the *train_model* function in RTextTools by calling the SVM algorithm from e1071.

NB classifier is a simple probabilistic classifier applied with Bayes Theorem with the assumption that all features are independent to each other [12]. Although this assumption may not be realistic in real-world situations, studies show that it outperforms other models for certain problems [13]. This model is built with the NB function in e1071.

ME is a classification method that generalises logistic regression to a multi-class problem and predicts the probabilities of different possible results. This algorithm is used as a contrast to NB, because of their opposite assumptions about the relationships between features. When the conditional independence assumption of NB fails in a certain case, ME might potentially yield a better result. ME algorithm is called from maxent [14] via the *train_model* function in RTextTools.

A classification tree is a predictive model using decision tree learning through a process of binary recursive partitioning. This is a popular approach for classification problems in data mining [11, chapter 8]. This is the foundation of other tree-based models that will be discussed later, including bagging, random forest and boosting. Tree algorithm is called from tree [15] via the *train_model* function in RTextTools.

Bagging, also called bootstrap aggregating is an ensemble algorithm used in machine learning to improve the accuracy and stability of an existing model. In this case a bagged classification tree is trained with ipred [16] and RTextTools.

Similar to bagged tree, RF is another ensemble learning algorithm. Instead of fully growing each binary tree, RF controls the number of features to search over in its pursuit of the best split for each tree. Using 200 trees, RF algorithm is called from randomForest [17] via the *train_model* function in RTextTools.

To reduce the bias of a predictive model against incorrectly classified data, boosting algorithms aim to build a strong

⁵<http://shiny.rstudio.com>

⁶<https://developers.google.com/chart>

TABLE I
OUTPUT OF SAMPLE ACCURACY GENERATED THROUGH 10-FOLD CROSS VALIDATION USING UNI-GRAM FEATURE.

Training Data Size	SVM	ME	Tree	Bagged Tree	Boosted Tree	RF	NN	NB
1000	60%	9.3%	56.8%	56.9%	84.4%	60.4%	50.3%	29.1%
2000	68.6%	7.2%	62.4%	65.3%	82.5%	68.6%	59.8%	40.3%
3000	73.5%	6.2%	65.6%	70.2%	84%	73%	66.4%	44.4%
4000	83.2%	5.2%	78.1%	80.6%	87.9%	83%	72.9%	43.7%
5000	84.5%	5.3%	78.3%	88.9%	88.5%	91.8%	78.8%	41.6%

TABLE II
OUTPUT OF SAMPLE ACCURACY GENERATED THROUGH 10-FOLD CROSS VALIDATION USING A TRAINING DATASET OF 5000 TWEETS.

N-grams	SVM	ME	Tree	Bagged Tree	Boosted Tree	RF	NN	NB
Uni-gram	84.5%	5.3%	78.3%	88.9%	88.5%	91.8%	78.8%	41.6%
Bi-gram	53.7%	2.3%	50.9%	53.7%	96.5%	52.9%	53.5%	41.3%
Tri-gram	44.2%	37.4%	43.9%	44.4%	99.6%	44.3%	44.2%	N/A

classifier on a weak one by iteratively adding weights to misclassified data [11, chapter 8]. Logit boosting algorithm from caTools [18] is employed here to build a boosted tree.

NN performs non-linear statistical modelling and is robust against noise and generalisation [19]. In this research we apply NN for text categorisation, which is seldom done. NN algorithm is called from nnet [20] via the *train_model* function in RTextTools.

C. Evaluation

Table I compares model performance under 40 conditions generated by different training data sizes of 1000, 2000, ..., 5000 tweets and algorithms using the uni-gram feature. Model performances enhance along with larger training data set, with the only exception of ME. The top four performing models from this evaluation are highlighted and chosen for classifying the data set - the best performing model is RF with nearly 92% accuracy, followed by Bagged Tree and Boosted Tree which have almost similar performances and the SVM which has almost 85% accuracy.

As shown in Table II, model performances are not improved by increasing N-gram features, with the exception of Boosted Tree reaching over 99% accuracy with tri-grams. The accuracy of ME using tri-gram feature, which is the other exception, and is nearly 20 times higher than using bi-gram. There is an ongoing debate on whether N-gram feature can effectively enhance model performance or not [12]. On one hand, uni-gram features cover more data where as bi-gram and tri-gram features better capture patterns of expressions. In this research, the uni-gram features are clearly the most effective and it is believed that the choice of N-gram features should depend on the specific corpus. When selecting well-performing models, the two exceptions above are regarded as special cases that cannot represent the general trend. Therefore, the four highlighted best classifiers in Table I, namely RF, Bagged Tree, Boosted Tree and SVM, are accepted for classifying the test data, the results of which will be compared visually in the web application.

Although NB is strongly recommended in [21] as a better choice for snippet sentiment tasks than SVM, it fails in this

case. Two factors might account for this result: first, the data sets used in other studies are usually a mixture of topics, such as movie reviews; second, the snippet sentiment analysis explores the short version of documents, which contain much less informal languages than tweets. Given the fact that this research focuses on the public sentiment towards one specific topic extracted from tweets, it is possible that so-called strong classifier such as NB could actually fail. NB using tri-gram feature is not available due to the lack of decisive phrases.

IV. TWITTER ANALYSIS OF SANDY HOOK ELEMENTARY SCHOOL SHOOTING

In this paper, we analyse the tweet corpus purchased from GNIP surrounding the Sandy Hook elementary school shootings of Friday, 14 December, 2012 at 0940 local (14:40 GMT) in Newtown, Connecticut, USA and visualise the results as a web application which is available at <http://www.gunsontwitter.com>.

A. Analysing the Tweets

The acquired Twitter corpus contains 700,437 tweets for a 40 day period from Friday, December 7 2012, 00:00:01 GMT until Tuesday, January 15 2012, 23:59:59 GMT. The relevant tweets were obtained from GNIP by applying a number of filter rules. The rules can incorporate a number of keywords (for example, Sandy Hook, Newtown and guns), phrases (for example, gun violence and firearm magazine), hashtags (for example, #SandyHook and #CTshooting) and reply/mention tags (for example, @sandyhook and @SHWeChooseLove) along with a variety of filters such as the country of origin of the tweet (country_code:us), language (lang:en) used for tweeting and a means to specify that the tweet is or is not a repost (-is:retweet) as no retweets were selected).

During pre-processing relevant information such as the tweets, date, time and location were extracted. A golden training data set of 5000 tweets were manually selected and classified on the basis of the following three sentiments towards the event

- 1) Pro-gun: favouring the freedom of gun ownership and against gun control. For example, "The only thing that stops




Fig. 2. Number of classified tweets from December 7 2012, 00:00:01 GMT to January 15 2012, 23:59:59 GMT

a bad guy with a gun, is a good guy with a gun” or “Gun control isn’t gonna help. God bless the 2nd amendment! #sorrynotsorry”.

- 2) Anti-gun: favouring gun control and express outrage and sympathy to the event. For example, “I cannot fathom how one could be so vile that they open fire on children or anyone. We need STRICT gun controls. #Newtown” or “BAN GUNS!!! Let our children be safe”.
- 3) Neutral: irrelevant in the context or neither convey a pro-gun or anti-gun sentiment. For example, “Obama is possibly visiting Newtown” or “Not sure if gunshot or firework”.

The four machine learning approaches that demonstrated best performance during evaluation in Section III, namely RF, Boosted Tree, Bagged Tree and SVM, were used to train models with the training data and then applied to classify the remaining tweets.

Figure 2 highlights the numbers of classified tweets using the four selected machine learning approaches. Pro-gun tweets account for about 10% of the data set regardless of approach chosen. The major difference between the approaches is in the number of neutral and anti-gun tweets. In the evaluation stage, a difference in accuracy by 10% was observed between the RF and SVM approaches. However, both these approaches assign over half of the tweets into the neutral class. The bagged tree model classifies nearly the same volume of tweets as neutral and anti-gun.

Figure 3 summarises the most frequently used hashtags and reply/mention tags in the data set. Although some of these tags do not convey obvious sentiments they act as indicators of emotions to some extent. For example #guncontrol and @BarackObama usually appear in anti-gun tweets while #NRA is more favoured in tweets that are more pro-gun and advocate gun ownership and gun rights. It would seem that there is more neutral sentiment conveying reply/mention tags (for example @justinbieber or @machinegunkelly).

B. Interactive Visualisation

An interactive web application (app), which is available on <http://www.gunsontwitter.com>, was developed to display multiple dimensions of the data. Data on population⁷ and

⁷<https://www.census.gov/popest/data/state/totals/2012/>

percentage of gun ownership (as of 2007)⁸ in each state were incorporated to add value to the visualisation. The following sections demonstrate the three visualisation techniques.

1) *Motion Chart*: Figure 4 is an example of the different visualisations possible using motion chart. For these figures the RF approach is selected. Figure 4a is a bubble view with six variables visualised simultaneously: X-axis for numbers of neutral tweets, Y-axis for public sentiment scores balanced by data volume and population ($PGPSS_3$), bubbles for the states, sizes of bubble for the populations of the states, colours of bubbles for the gun ownership as percentage of population and the slider at the bottom for date. In this chart for December 7 2012 (one week before the shooting event), larger the population of the state more pro-gun is the sentiment for gun use. In contrast, states that have a large volume of registered gun owners show little pro-gun sentiment.


A trail feature can be set by the user for viewing the variation of any variable over a time period. For example, refer Figure 4b. Texas is selected, and an overview of its variation during the 40-day period is seen on the chart. The highest $PGPSS_3$ appears when the number of neutral tweets increased to its top level.

Figure 4c shows the embedded bar chart view provided by the motion chart. Five groups of statistics are presented: numbers of pro-gun tweets on X-axis, volumes of tweet on Y-axis, states as bars, populations in the colours of bars and dates on the slider. In this example, the results of tweets on December 16 2012 are presented. The majority of tweets come from states with large population except Connecticut standing out in the middle of X-axis. This is probably because the shooting happened there and people were more concerned.


Figure 4d is the line graph view of the motion chart. Four variables are available on this view: time series on X-axis, numbers of anti-gun tweets on Y-axis, states shown in lines and public sentiment scores balanced by volumes of data ($PGPSS_2$) represented as the colour of the lines. This graph compares the statistics at the state-level and cannot provide a holistic view at the country-level. Therefore, a different line graph is generated as the second visualisation technique to analyse the results of U.S. at the country-level.

2) *Line Graph*: Figure 5 presents the visualisation using line graph. These results are presented for the Bagged Tree model. The numbers of tweets in the three sentiment classes can be explored either on a daily (Figure 5a) or an hourly (Figure 5b) basis. Zooming options are realised through a series of buttons at the top left corner, from one hour to the maximum time period. The time frame at the bottom is designed to customise the view a certain period. By default the line graph displays the results for the 40-day period. These graphs highlight that during the one-week period past the shooting, the number of tweets increase considerably regardless of the sentiment (pro-gun or anti-gun) with peaks observed on December 14 2012. An interesting fact is that although

⁸<http://usliberals.about.com/od/Election2012Factors/a/Gun-Owners-As-Percentage-Of-Each-States-Population.htm>



(a) Hashtags



(b) Reply/Mention Tags

Fig. 3. Top 20 tags frequently used in tweets from December 7 2012, 00:00:01 GMT to January 15 2013, 23:59:59 GMT

anti-gun sentiment was normally stronger than the pro-gun sentiment, on December 19 2012 the anti-gun sentiment is as low as the pro-gun sentiment. This five-day period (from December 14 to 19) could be important for using public opinion as an instrument for facilitating policy making.

3) *Geographic Map*: The coordinates of each tweet (latitude and longitude) were used to categorise the tweets from each state. This was facilitated using four R packages, namely *regos* [22], *sp* [23], *maps* [24] and *maptools* [25]. The coordinate information was used to generate a geographic visualisation as shown in Figure 6 and Figure 7.

Figure 6 presents three maps based on $PGPSS_1$, $PGPSS_2$ and $PGPSS_3$ for the 40-day period. $PGPSS_3$ is the best estimation since it is corrected for both volume of data and population. $PGPSS_1$ may not be a good indicator since it only takes the ratio of pro-gun and anti-gun tweets into account. In addition, the percentage of gun ownership appears in a floating box when the cursor hovers over a state. By default the geographic map displays $PGPSS_3$ over the whole 40-day period until the user chooses dates of interests through the date picker. For example, refer Figure 7. By comparing Figure 6c and Figure 7c, the three states that have strong pro-gun sentiment are California, Texas and New York.

V. DISCUSSION


The Sandy Hook school massacre was a deeply distressing and tragic event. It raised issues related to gun ownership and mental health. The mental health issue is not addressed by our data but we note that others who have carefully considered the issue [26] conclude that “notions of mental illness that emerge in relation to mass shootings frequently reflect larger cultural stereotypes and anxieties...”. The focus thus logically returns to gun ownership and access. High profile events typically evoke a huge and often opinionated social media response [27], [28],

[29]. In instigating the analysis our prior assumption was that there would be a peak of anti-gun sentiment expressed in the hours immediately following the shooting. This did indeed turn out to be the case. We also assumed that after a period of time there may be a pro-gun response; that this response may come in reply to calls for greater gun control and would be seen by those involved as a reassertion of their second amendment rights. Our assumption was that this response may have occurred after about 7 to 10 days and if so we were interested in delineating the period between the anti-gun and pro-gun spikes as this perhaps defined a period during which preventative political action could be initiated.

Surprisingly, at least to us, the pro-gun spike in sentiment was not delayed but actually occurred on the day of the killings. We do not assume that this was a heartless response to a truly shocking and distressing event. It is quite possible that some of those expressing pro-gun sentiments were unaware that there had been fatalities at the time they tweeted. For good operational and practical reasons the fact that there had been so many deaths did not start to appear in the media until quite some hours after the event.

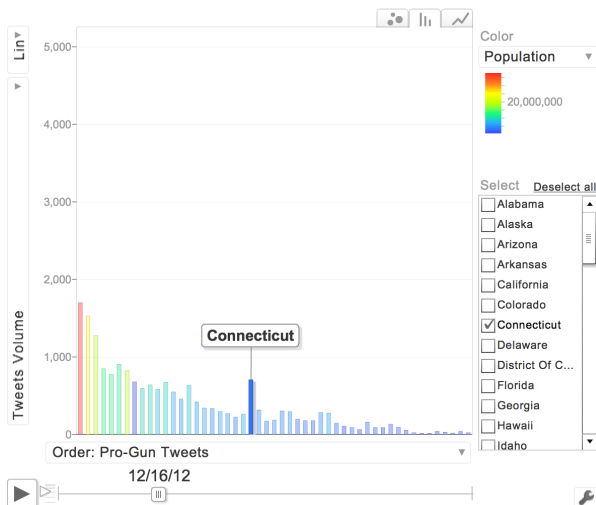
It is also interesting to note that whilst anti-gun sentiment peaked and then quickly fell to pre-shooting levels, pro-gun sentiment ran at an elevated level for some time. Again we think this probably reflects a genuine and deeply ingrained difference in opinion between individuals about the role of guns in causing or preventing such mass shooting events.

Indeed we think it is likely that pro-gun tweeters on the day of the killings did so because they genuinely believe that the event could have been avoided had members of staff been armed. A full discussion of the evidential basis of any such presumed belief or indeed of its more general moral, social and educational consequences is beyond the scope of this paper. State level analysis showed considerable variation in




(a) Bubble view

(b) Trail feature of the bubble view



(c) Bar chart view



(d) Line graph view

Fig. 4. Motion chart of the interactive web application

response. The highest gun owning states were not the most pro-gun in this situation. This may be because reasons for gun ownership vary considerably. Some traditional hunting states with high level of gun ownership demonstrated relatively low levels of pro-gun sentiment. California, Texas and New York demonstrated high levels of pro-gun sentiment. A full exploration of this interesting observation is beyond the scope of the data at the heart of this paper.


VI. RELATED WORK

Sentiment analysis of social media is an important avenue of research in the area of natural language processing. Previous research includes document-level analysis of product reviews [30] to term-level inspection on the polarity of words [31] [32]. Since tweets are character limited sentence-level classifications are frequently adopted to extract public sentiment. Classifying tweets is challenging due to the unique nature of micro-


blogging with its frequent use of informal and colloquial language including slang and emoticons. Although there are different approaches for classifying tweets, there is no consensus on the best solution. In natural language processing of tweets, engineering linguistic features and automating text categorisation are two important tasks.

A. Linguistic Features

Twitter sentiment has been explored from the traditional linguistic perspectives and extracting N-gram features is a baseline method among researchers [6]. Uni-gram is the most common feature used and is reported to result in good predictive models with accuracy around 70%-80% [33] [34]. Bi-gram and tri-gram features are claimed to be effective in improving model performances. However, research suggests that uni-grams are better when building classifiers [12].



(a) Daily Analysis



(b) Hourly Analysis

Fig. 5. Line graph of the interactive web application




Fig. 6. Geographic map with customised PGPS for December 7 2012 to January 15 2013

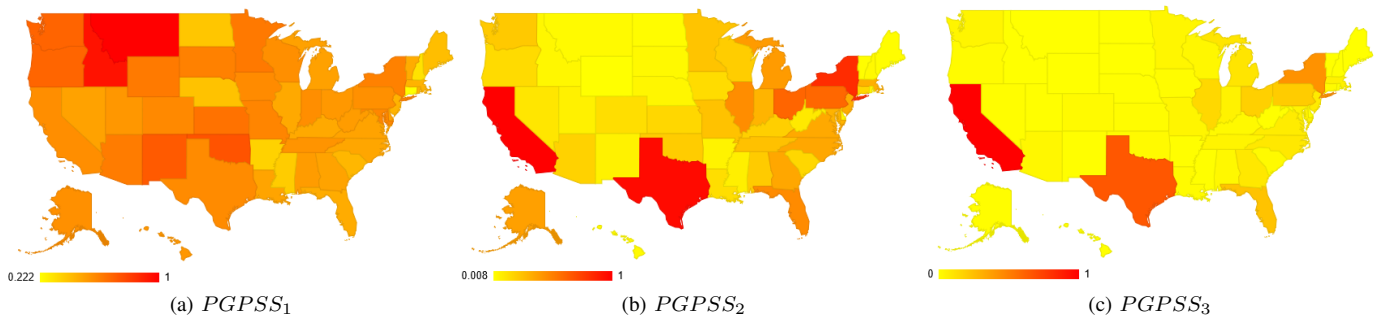


Fig. 7. Geographic map with customised PGPS for December 13 2012 to December 15 2012

Other features such as Part-Of-Speech (POS) tags [7] and the position of words in text [12] are strong indicators of emotional text [13]. However, experiments show that POS tags may not be useful for sentiment analysis of micro-blogs [6]. Hash-tags and emoticons are typical contents in Twitter that indicate latent sentiment which are made use of in the process of automatically generating training data [35] [36].

B. Automatic Text Categorisation

Algorithms for automated classification include dictionary-based and machine learning approaches. Dictionary-based approaches [37] [38] compare tweets against lexicons of existing dictionaries. By counting frequencies of the words matched with the chosen dictionary, a tweet is assigned a sentiment based on the prevalent emotion in the text. This approach is disadvantageous in that the acquired lexicon is domain independent and therefore it is impossible for a dictionary-based algorithm to capture different contexts. Another drawback of this approach is that it has to ignore the informal language used in tweets, which could probably stand for strong emotions. For example, words with repeated letters such as ‘nooooooooo’ convey a stronger emotion than ‘no’. To overcome these disadvantages, machine learning based approaches provide the possibility of enhanced utility.

Machine learning approaches can be incorporated in building classifiers and can produce accuracy comparable to human experts [39]. However, the best classifier depends on the specific case study it is employed for; SVM and NB are popularly used owing to their average accuracy of above 80% regardless of features selected [7]. Other algorithms such as Conditional Random Field (CRF) and Maximum Entropy (ME) are found to be less effective for tweets analysis [13] [7].

In addition to the aforementioned standard machine learning techniques, adaptive classifiers have been built to improve performance. For example, the NBSVM, an SVM classifier trained with NB features, has better performance than SVM [21]. Similarly, an interpolation method [34] can be employed to modify the NB classifier with a generative model of words given semantic concepts. Hybrid classification has been proposed as an effective technique to combine a rule-based classification and machine learning approach [40].

C. Challenges

Although functional techniques have been developed in the area of using machine learning for sentiment analysis, there continues to be three challenges that we seek to address in this paper. Firstly, there is a lack of evaluation of various machine learning approaches in the same Twitter analysis context. For example, the so-called best classifiers, which were trained for categorising positive and negative text, demonstrated poor performance with three classes (positive, negative and neutral) [7]. Meanwhile, approaches such as Neural Network (NN) and Boosting in sentiment analysis are scarcely investigated. Hence, we believe it is essential to evaluate different ap-

proaches in the same context in order to identify the best approach in any given context.

Secondly, an important social issue, namely gun violence, has not been investigated for understanding public sentiment. Twitter analysis has been employed for understanding complex phenomena in the areas of politics [37], disaster management [33] and public health [41]. Gun violence is an important and challenging topic in public health and understanding the temporal and geographic patterning of pro-gun and anti-gun sentiments is important in terms of framing and timing political initiatives aimed at reducing gun related deaths. This is particularly true in terms of mass shootings and in this paper, we seek to use machine learning based Twitter analysis to explore pro-gun and anti-gun sentiments in relation to the Sandy Hook School shootings.

Thirdly the use of machine learning techniques for the analysis of large data sets is in general still confined to the academic discipline of Computer Science (CS). In seeking to bridge the disciplines of Computer Science and Public Health in producing this paper we are seeking to further the prospects of such trans-disciplinary research by promoting the benefits that machine learning can offer.

Hence, in this study, the above three challenges related to the evaluation of machine learning techniques, the application of Twitter analysis to gun violence and the accessibility of machine learning techniques for non-CS experts are addressed through the Sandy Hook Elementary School shootings case study. Tweets expressed public sentiment in terms of pro-gun or anti-gun attitudes were examined over forty days. These sentiments were extracted from tweets using machine learning techniques and visualised.

VII. CONCLUSIONS

The public health implications are firstly methodological. The disciplines of Computer Science and Public Health are brought together on the important and emotive issue of reactions to a mass school shooting which is a useful exercise that yields multidisciplinary dividends.

Furthermore in this paper, we have evaluated a number of machine learning approaches and identified those most suitable to classifying public sentiment towards gun violence in light of the Sandy Hook school shooting. We have shown that it is possible to analyse a large body of social media data using machine learning in a reliable and replicable way by employing a methodology to collect, train and classify tweets. Public sentiment is captured for different time scales and by state can be interactively visualised on an online tool available at <http://www.gunsontwitter.com>.

Sentiment analysis shows a peak of anti-gun feeling on the day of the Sandy Hook school shooting which quickly falls to pre-event levels. More surprisingly the analysis shows a peak of pro-gun sentiment on the day of the shooting that is sustained at an elevated level for a number of days. These findings suggest that a discernible minority of the US population see high levels of public gun ownership as part of the solution rather than part of problem. This is in spite of the

objective evidence that the US with its generally liberal gun ownership rules suffers a uniquely high incidence of school shootings [42] and that high level of public firearms ownership increase firearm related deaths both of the public [43] and of law enforcement officers [44].

More work is required to understand the complexities of public opinion on this important issue but social media analysis appears to have a useful contribution to make. In the future, we aim to apply machine learning techniques to additional gun violence related events. This will enhance our knowledge of public reaction and delineate a time period during which preventative political action can be instigated.

REFERENCES

- [1] S. Sedensky, "Report of the state's attorney for the judicial district of danbury on the shootings at sandy hook elementary school and 36 yogananda street, newtown, connecticut on december 14, 2012, state of connecticut," November 2013.
- [2] J. M. Shultz, G. W. Muschert, A. Dingwall, and A. M. Cohen, "The sandy hook elementary school shooting as tipping point," *Disaster Health*, vol. 1, no. 2, pp. 65–73, 2013.
- [3] S. Li, G. Zhou, Z. Wang, S. Y. M. Lee, and R. Wang, "Imbalanced sentiment classification," in *Proceedings of the 20th ACM International Conference on Information and Knowledge Management*, pp. 2469–2472, 2011.
- [4] P. Zhang, "Model selection via multifold cross validation," *The Annals of Statistics*, pp. 299–313, 1993.
- [5] M. Gesmann and D. de Castillo, "googlevis: Interface between r and the google visualisation api," *The R Journal*, vol. 3, no. 2, pp. 40–44, 2011.
- [6] E. Kouloumpis, T. Wilson, and J. Moore, *Twitter Sentiment Analysis: The Good the Bad and the OMG!*, pp. 538–541. AAAI Press, 2011.
- [7] A. Go, R. Bhayani, and L. Huang, "Twitter sentiment classification using distant supervision," *CS224N Project Report, Stanford*, pp. 1–12, 2009.
- [8] V. D. Nguyen, B. Varghese, and A. Barker, "The royal birth of 2013: Analysing and visualising public sentiment in the uk using twitter," in *Proceedings of the IEEE International Conference on Big Data*, pp. 46–54, 2013.
- [9] E. Dimitriadou, K. Hornik, F. Leisch, D. Meyer, A. Weingessel, and M. F. Leisch, "The e1071 package," *Department of Statistics (e1071)*, TU Wien., 2006.
- [10] T. P. Jurka, L. Collingwood, A. E. Boydston, E. Grossman, and W. van Atteveldt, "Rtexttools: Automatic text classification via supervised learning," *R package version 1.3.9.*, 2012.
- [11] G. James, D. Witten, T. Hastie, and R. Tibshirani, *An introduction to statistical learning*. Springer, 2013.
- [12] B. Pang, L. Lee, and S. Vaithyanathan, "Thumbs up?: sentiment classification using machine learning techniques," in *Proceedings of the ACL-02 conference on Empirical methods in natural language processing-Volume 10*, pp. 79–86, Association for Computational Linguistics, 2002.
- [13] A. Pak and P. Paroubek, "Twitter as a corpus for sentiment analysis and opinion mining," in *Proceedings of the Seventh conference on International Language Resources and Evaluation*, 2010.
- [14] T. P. Jurka, "maxent: An r package for low-memory multinomial logistic regression with support for semi-automated text classification," *The R Journal*, vol. 4, no. 1, pp. 56–59, 2012.
- [15] B. Ripley, "Classification and regression trees," *R package version 1.0-35*, 2005.
- [16] A. Peters and T. Hothorn, "ipred: improved predictors," *R package version 0.9-3*, 2013.
- [17] A. Liaw and M. Wiener, "Classification and regression by randomforest," *R News*, vol. 2, no. 3, pp. 18–22, 2002.
- [18] J. Tuszynski, "catools: Tools: moving window statistics, gif, base64, roc auc, etc.," *R package version 1.17*, 2014.
- [19] I. Basheer and M. Hajmeer, "Artificial neural networks: fundamentals, computing, design, and application," *Journal of microbiological methods*, vol. 43, no. 1, pp. 3–31, 2000.
- [20] W. N. Venables and B. D. Ripley, *Modern Applied Statistics with S*. New York: Springer, fourth ed., 2002.
- [21] S. Wang and C. D. Manning, "Baselines and bigrams: Simple, good sentiment and topic classification," in *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Short Papers-Volume 2*, pp. 90–94, Association for Computational Linguistics, 2012.
- [22] R. Bivand and C. Rundel, "rgeos: Interface to geometry engine - open source (geos)," *R package version 0.3-6*, 2014.
- [23] E. J. Pebesma and R. S. Bivand, "Classes and methods for spatial data in R," *R News*, vol. 5, pp. 9–13, November 2005.
- [24] R. Brownrigg, "maps: Draw geographical maps," *R package version 2.3-7*, 2014.
- [25] R. Bivand and N. Lewin-Koh, "mapproj: Tools for reading and handling spatial objects," *R package version 0.8-30*, 2014.
- [26] J. M. Metzler and K. T. MacLeish, "Mental illness, mass shootings, and the politics of american firearms," *American Journal of Public Health*, vol. 105, no. 2, pp. 240–249, 2015.
- [27] P. Burnap, M. L. Williams, L. Sloan, O. Rana, W. Housley, A. Edwards, V. Knight, R. Procter, and A. Voss, "Tweeting the terror: modelling the social media reaction to the woolwich terrorist attack," *Social Network Analysis and Mining*, vol. 4, no. 1, pp. 1–14, 2014.
- [28] G. Lotan, E. Graeff, M. Ananny, D. Gaffney, I. Pearce, and danah boyd, "The arab spring— the revolutions were tweeted: Information flows during the 2011 tunisian and egyptian revolutions," *International Journal of Communication*, vol. 5, no. 0, 2011.
- [29] R. Procter, F. Vis, and A. Voss, "Reading the riots on twitter: methodological innovation for the analysis of big data," *International Journal of Social Research Methodology*, vol. 16, no. 3, pp. 197–214, 2013.
- [30] B. Pang and L. Lee, "Opinion mining and sentiment analysis," *Foundations and trends in information retrieval*, vol. 2, pp. 1–135, 2008.
- [31] A. Esuli and F. Sebastiani, "Sentiwordnet: A publicly available lexical resource for opinion mining," in *Proceedings of the International Conference on Language Resources and Evaluation*, vol. 6, pp. 417–422, 2006.
- [32] S. Mohammad, C. Dunne, and B. Dorr, "Generating high-coverage semantic orientation lexicons from overtly marked words and a thesaurus," in *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing: Volume 2 - Volume 2*, pp. 599–608, 2009.
- [33] 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*, pp. 851–860, 2010.
- [34] H. Saif, Y. He, and H. Alani, "Semantic sentiment analysis of twitter," in *Proceedings of the 11th International Conference on The Semantic Web - Volume Part I*, pp. 508–524, 2012.
- [35] A. Bifet and E. Frank, "Sentiment knowledge discovery in twitter streaming data," in *Proceedings of the 13th international conference on Discovery Science*, pp. 1–15, 2010.
- [36] D. Davidov, O. Tsur, and A. Rappoport, "Enhanced sentiment learning using twitter hashtags and smileys," in *Proceedings of the 23rd International Conference on Computational Linguistics*, pp. 241–249, 2010.
- [37] A. Tumasjan, T. O. Sprenger, P. G. Sandner, and I. M. Welpe, "Election forecasts with twitter: How 140 characters reflect the political landscape," *Social Science Computer Review*, p. 0894439310386557, 2010.
- [38] S. Asur and B. A. Huberman, "Predicting the future with social media," in *Proceedings of the 2010 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology*, pp. 492–499, 2010.
- [39] F. Sebastiani, "Machine learning in automated text categorization," *ACM computing surveys*, vol. 34, no. 1, pp. 1–47, 2002.
- [40] R. Prabowo and M. Thelwall, "Sentiment analysis: A combined approach," *Journal of Informetrics*, vol. 3, no. 2, pp. 143–157, 2009.
- [41] M. J. Paul and M. Dredze, "You are what you tweet: Analyzing twitter for public health," in *Proceedings of the 5th International AAAI Conference on Weblogs and Social Media*, pp. 265–272, 2011.
- [42] N. Böckler, T. Seeger, P. Sitzer, and W. Heitmeyer, *School Shootings: International Research, Case Studies, and Concepts for Prevention*, ch. School Shootings: Conceptual Framework and International Empirical Trends, pp. 1–24. 2013.
- [43] M. Siegl, Y. Negussu, S. Vanhoi, J. Pleskunas, C. S. Ross, and C. King 3rd, "The relationship between gun ownership and stranger and nonstranger firearm homicide rates in the united states, 1981–2010," *American Journal of Public Health*, vol. 104, no. 10, pp. 1912–1919, 2014.
- [44] D. I. Swedler, M. M. Simmons, F. Dominici, and D. Hemenway, "Firearm prevalence and homicides of law enforcement officers in the united states," *American Journal of Public Health*, vol. 105, no. 10, pp. 2042–2048, 2015.