

Application of Fuzzy Semantic Similarity Measures to Event Detection Within Tweets

Keeley Crockett, Naeemeh Adel, James O'Shea, Alan Crispin,
Intelligent Systems Group, School of Computing,
Mathematics and Digital Technology,
Manchester Metropolitan University, Chester Street,
Manchester, M1 5GD, UK
K.Crockett@mmu.ac.uk

David Chandran
Institute of Psychiatry, Psychology & Neuroscience,
Kings College London, 16 De Crespigny Park,
London, SE5 8AF, UK
João Paulo Carvalho
INESC-ID / Instituto Superior Técnico,
Universidade de Lisboa, Portugal

Abstract— This paper examines the suitability of applying fuzzy semantic similarity measures (FSSM) to the task of detecting potential future events through the use of a group of prototypical event tweets. FSSM are ideal measures to be used to analyse the semantic textual content of tweets due to the ability to deal equally with not only nouns, verbs, adjectives and adverbs, but also perception based fuzzy words. The proposed methodology first creates a set of prototypical event related tweets and a control group of tweets from a data source, then calculates the semantic similarity against an event dataset compiled from tweets issued during the 2011 London riots. The dataset of tweets contained a proportion of tweets that the Guardian Newspaper publically released that were attributed to 200 influential Twitter users during the actual riot. The effects of changing the semantic similarity threshold are investigated in order to evaluate if Twitter tweets can be used in conjunction with fuzzy short text similarity measures and prototypical event related tweets to determine if an event is more likely to occur. By looking at the increase in frequency of tweets in the dataset, over a certain similarity threshold when matched with prototypical event tweets about riots, the results have shown that a potential future event can be detected.

Keywords- fuzzy semantic similarity measures, Twitter, semantic analysis

I. INTRODUCTION

Twitter continues to be key player in the social media market with reporting on average 317 million monthly active users in late 2016 [1] and has the opportunity to be a rich source of information for organisations. However, the most difficult challenge lies in extracting meaning from the unstructured and inherently noisy textual tweets. Typically, to extract useful information from tweets i.e. sentiment analysis [2,3], tweets undergo significant pre-processing that can include removal of all URLs, correction of spelling errors, tagging of named entities, removal of stop words, punctuation etc., acronym look up and even the removal of hashtags. The cleaned tweet can, in cases, project a different semantic meaning than what was intended in the context of the analysis being conducted.

Tweets are known as a mechanism for spreading news information fast. For example, tweets were influential during the Arab Spring uprising in 2010-2011 [3], allowing large groups of people to communicate quickly and organise

protest rallies against regimes. Twitter also, perhaps more importantly, allowed information in the form of photographs and videos to be broadcast by members of the public. Munroe [4] suggested that people's communications over the Internet, especially Twitter could overtake an initial earthquake. This is because seismic waves travel a lot more slowly than data traveling along fiber optical cables. This effect has occurred a couple of times, more recently on 23rd August 2011 where a 5.9 magnitude earthquake struck close to Richmond, Virginia. The effects were initially felt in Washington D.C. where the initial tweets were posted, and various people reported having read them in New York City before the earthquake reached them [5,6].

Event detection and user profiling using Twitter is a predominant research area [2..13]. An event can be defined simply as an occasion of importance that happens at a given time. Sakaki et al [8] states that events have three key properties: they are often large scale, have an influence on an individual person's daily lifestyle and can have both spatial and temporal locations. Early research in [8] used support vector machines to classify a tweet into positive and negative classes in relation to the event being predicted. This classification was used to semantically analyse tweets and in conjunction with spatial estimation was incorporated into an earthquake reporting system. The authors [8] identified that the search query or term(s) used for classification are of vital importance and improvement in recall heavily relied on this factor. Pavlyshenko [9], used frequent sets, association rules and formal concept analysis to build semantic concepts between individuals, had more success detecting events that had fewer random factors. Arias et al [10] used summary decision trees and support vector machines to improve the power of Twitter forecasting models in predicting the stock market and box office revenue trends. Rui et al [11] proposed a Twitter based event detection and analysis system for crime and disaster related events. Ribeiro et al [12] proposed a method to identify traffic events and conditions using Twitter and report them in real time (with 50% to 90% accuracy). More recently, Polhl and Bouchachia [13] conducted a review of how social media networks were used to disseminate information during a crisis i.e. a police emergency. To the authors' knowledge, the incorporation of

semantic similarity as an additional dimension to the models produced was not considered.

Semantics are concerned with the literal meaning of morphemes, words, phrases and sentences and the way that they are combined. Semantic similarity is therefore, a complex concept with a long history in cognitive psychology and linguistics [14,15], which can analyse the deep semantic structure of a short text to convey meaning. An operational definition often used in studies of semantic similarity is “How close do these two sentences come to meaning the same thing?” [16] and getting a machine to be able to answer the question in a similar way to a human being is particularly challenging.

Due to the inherent natural language of tweets, fuzzy sentence similarity measures (FSSM) are of particular interest in this work. FSSM are algorithms that are able to compare two or more short texts which contain human perception based words and return a numeric measure of similarity of meaning between them. FAST [15] (Fuzzy Algorithm for Similarity Testing) is an ontology based similarity measure that uses concepts of type 1 fuzzy sets to model relationships between categories of human perception based words (fuzzy words). Previous work has shown that FAST gives higher correlations with human ratings of similarity than leading other measures [17], which tend to ignore fuzzy properties of words when measuring similarity. To the knowledge of the authors, none of the work on event detection to date has measured the semantic similarity of the tweets using fuzzy short text similarity measures. A brief review of FAST can be found in Section II.

The aim of the research presented in this paper is to see if groupings of prototypical tweets about a potential event, i.e. a riot, can be used in conjunction with fuzzy short text similarity measures to detect where an event is more likely to occur. For the purpose of this work the chosen event is the London riots [18,19] which took place between the 6th and 10th August 2011, where the UK experienced riots at a level not seen since the eighties. Following the riots, not only did the UK Government announce a public enquiry, the Guardian newspaper began its own analysis which included an examination of the role of social media to try and establish whether Facebook and/or Twitter actually incited the riots [20,22]. This analysis involved examination of 2.57m tweets and concluded that tweets during the period were mainly used as a reaction mechanism. The research in our paper addresses the question that if the tweets of the perpetrators were known days before the actual riots and provided evidence that the riots were incited in some way, could these individuals have been brought to justice sooner? Also, could we potentially avoid or scale down the riots themselves? This would have positive effects for society, including reduced insurance claims. If we can detect that a potentially dangerous or criminal event is about to occur and identify who is initiating it and where it is likely to occur, we can then put measures in place to either stop it happening or reduce the consequences.

This paper is organised as follows; Section II provides an overview of fuzzy short text semantic similarity measures. Section III describes how the London Riot Data set was created and how the data was sampled for this work. Section IV presents a methodology for application of fuzzy semantic similarity measures in detecting potential events using semantic analysis experimental results and accompanying discussion are covered in Section V and finally Section VI presents the conclusions and future directions.

II. OVERVIEW FUZZY SHORT TEXT SEMANTIC SIMILARITY MEASURES

FAST (Fuzzy Algorithm for Similarity Testing) [17] was developed to enable new human perception properties to be taken into consideration when short texts were analysed by a machine to determine their syntactic and semantic similarity. FAST was inspired by STASIS, a short text semantic similarity measure developed by Li et al [23], from which FAST adopted the path length and depth of words relative to their position in a set of fuzzy ontologies with the information content of individual words being derived from a corpus. These were used to form semantic vectors and were then combined with word order vectors (from the word order in each short text) to determine the semantic similarity [17]. FAST identified fuzzy words within a short text and calculated the effect such words would have on the overall similarity. Experimental results showed that FAST gave an improved correlation between the similarity measure and human ratings [18, 24] compared with traditional measures.

Essentially, FAST works through first applying a word similarity measure to every possible pair of words in a short text and using corpus statistics to determine the overall semantic similarity between two short texts. The key stages of the FAST measure are:

- 1) Tokenize every word in the two short texts. For example, a cleaned tweet (method outlined in section IV) (“UKRiots Those convicted include a primary school teacher a lifeguard a man who works for homeless charity and an 11 year old from Essex”) is sorted into a list [“UKRiots”, “Those”, “convicted”, “include”, “a”, “primary”, “school”, “teacher”, “a”, “lifeguard”, “a”, “man”, “who”, “works”, “for”, “homeless”, “charity”, “and”, “an”, “11”, “year”, “old”, “from”, “Essex”]
- 2) Pair every combination of tokenized words. A Bag of Words was [23] created as the union between all words within the two short texts which similarity is being measured.
- 3) Determine the similarity of each word pair. If the word pair comprises of only fuzzy words i.e. [young, old] then use fuzzy category based ontologies to determine path length else use WordNet [25] as the semantic knowledge base to calculate path length.
- 4) If the word pair contained non-fuzzy words i.e. [teacher, man], determine effect of associated fuzzy words i.e. “small man”.

- 5) Apply sentence similarity measure using word similarities from different word pair combinations from 3) and 4).

Developed from an established traditional sentence similarity measure known as STASIS [24], FAST also incorporates an empirically determined semantic threshold, α , which was used to filter out word pairs with very low similarity scores. Li et al [24] justified the use of a semantic threshold, particularly when short texts were very short in length. The work also determined that function words (words that express grammatical relationships with other words within a short text [24] i.e. 'do'), also carried syntactic information and were to be included in the semantic similarity measurement. Typically, tweets are short in length (i.e. 140 characters per tweet or less excluding multimedia). In this work, the application of FAST to determine the semantic similarity to tweets will require the empirical evaluation of a suitable semantic threshold. FAST is fully automatic without requiring the users' intervention and readily adaptable across the range of potential application domains. For the purpose of this research FAST will be used to measure the similarity of Twitter tweets to a set of prototypical tweets in an attempt to detect a potential future event.

III. CREATING THE LONDON RIOT DATA SET

In order to evaluate the use of FSSM in its suitability to detect possible future events from tweets, it was necessary to construct a dataset to investigate whether or not an event could be predicted. Given the number of tweets generated on a daily basis, it was essential that the dataset contained a balanced proportion of tweets that concerned a particular event. The event selected for this study was the London Riots which occurred between the 6th and 10th August 2011. The riots were seen to be triggered by the shooting of 29-year-old father of four Mark Duggan by the police. The Guardian Newspaper publically released some Twitter data that included a list of 200 influential Twitter users based on re-tweets during the riot period [21]. It also included a list of the most popular Hashtags - relevant keywords, acronyms or phases in order to allow the tweet to be categorised. The dataset, known as the London-Riot dataset was initially populated with tweets from users identified using the Guardian data. The dataset was then expanded by selecting users, which appeared using the Twitter REST API public feed. For each user, tweets were recorded which were created up to and after 1st August 2011 at midnight, or up to the 3,200 tweet limit from the REST API statuses/user_timeline limitation (if the user had posted more than 3,200 tweets since 1st August 2011). A total of 9,913,397 tweets were collected from 8,819 Twitter users.

Due to the time taken to process this quantity of data, using available equipment this dataset was further reduced in size. A total of 1,132,938 individual tweets were extracted between 1st August 2011 00:00:00 and 31st August 2011 23:59:59 to create a new dataset which will be referred to in

this work as the Twitter Riot dataset. The quantity of Guardian riot tweets which appeared in the results was 17,795 tweets – a total of 4.6% of the tweets collected were sourced from users listed in the Guardian data. Samples of tweets from an 11 day period were then extracted to test groupings of prototypical tweets. Details are provided in Section IV.

IV. DETECTING EVENTS USING FUZZY SHORT TEXT SEMANTIC SIMILARITY MEASURES

This section describes a study that was conducted to test the following hypothesis:

H1: Can Twitter tweets be used in conjunction with fuzzy short text similarity measures and prototypical event related tweets to determine if an event is more likely to occur.

The section first outlines the overall methodology and then describes how the prototypical tweets for an event were sampled and a control group formulated.

A. Methodology

Let Event dataset be a generic name used to define a set of tweets that is to be used to investigate the likelihood of an event occurring that is time stamped. For this work, 2200 tweets were randomly sampled without replacement from the Twitter Riot Dataset; 200 tweets were randomly sampled from the period 1st to 11th August to investigate before, during and after the riots started on the 6th April 2011. The methodology for using fuzzy short text semantic similarity measures to detect potential events is defined as follows:

1. Select a series $\{1..k\}$ where $1 \leq k \leq m$ of prototypical tweets, T concerning an event, where m is the maximum number of associated prototypical tweets and is empirically defined. In this work $k = 7$. Each tweet, t is between 1 and 25 words in length.
2. For all tweets in the Event dataset $\{1..n\}$ where n is the number of tweets, calculate the fuzzy semantic similarity, S_i between every tweet, t_n , per day stored from the *start-date* to the *end-date* of the month of the event and prototypical tweet km using the pre-selected FSSM.
3. Using the short text semantic similarity measure, S_i , plot a graph showing the following:
 - a. The total number of tweets stored for each day between the start-date and end-date.
 - b. The number of tweets per day where the similarity S_i of t_i is greater than a given semantic threshold, α , where $0.5 \leq \alpha \leq 0.7$ in 0.05 increments. A semantic threshold, α , of 0.5 was chosen as an initial starting point, as this will result in matches which have a moderate to high similarity with the comparative short text.
4. Identify if there is sufficient increase in the frequency of tweets in the dataset, over a specific similarity threshold during the event. This sufficient increase will be

identified if the comparison result, from one day to another, is greater than 0.5%.

No cleaning of the tweets took place prior to running through FAST. FAST removes symbols, such as (\$%^|&* etc.) from the short texts, leaving only letters and numbers. For example, the anonymized tweet “\@XXX77 Tottenham has a notorious past for riots, Im sure it aint the last time, :0(“ becomes “XXX77 Tottenham has a notorious past for riots, Im sure it aint the last time 0”. This ensures word order, path length and the inclusion of function words is maintained – all required to determine the similarity. Hashtag words are also left in place with “#Riots:” becoming “Riots”.

B. Selection of Prototypical Tweets

In order to evaluate if FAST could be used to detect if a potential event was more likely to occur, two groups of tweets were selected. The first group contained 7 tweets that were related to the type of event that was to be potentially detected – in this case riots. Tweets ID’s 1 to 7 (Table I) were randomly selected from a study on “*Twitter, Information Sharing and the London Riots*” [22] which analyzed 600,000 tweets and retweets about the London riots to investigate whether Twitter was used as a tool to promote illegal group actions. The second group, known as the control group (Tweet ID’s 8 to 14) contained a further 7 tweets which were randomly sampled from top tweets of 2011 [26]. The 14 prototypical and control group tweets can be seen in Table I.

C. Experimental Methodology

All 2200 tweets were ran against the 7 prototypical event tweets and the 7 control tweets shown in Table I, for each of the semantic thresholds and the similarity of each was recorded. In order to identify if the semantic similarity of the tweets indicated if an event was likely to occur, each day’s tweets which matched the cumulative prototypical tweets over a specific similarity threshold, α had to be scaled as a % of that day’s tweets. The relative number of tweets on that given day defined as

$$\%relative\ tweets = \frac{tweets\ measuring\ above\ \alpha}{tweets\ processed\ for\ given\ date} * 100 \quad (1)$$

Where α is in {0.5, 0.55, 0.6, 0.65, 0.7}.

V. RESULTS AND DISCUSSION

A. Results

The results were grouped into the date that each tweet was posted, in order for a day-by-day comparison. Tables II and III show the % relative tweets for both the riot prototypical tweets and the control tweets over the eleven-day period, along with the semantic threshold used in each experiment. Higher semantic similarity thresholds did not yield an increase, emphasising the need that prototypical event tweets need to be more generalised to an event type and not a specific occurrence of an event. From Table II, it can be seen that

during the days preceding the first riot on the 6th August and in the days afterwards when the rioting spread to further cities in the north (6th to 11th) there was a higher number of tweets which matched the prototypical tweets with high semantic similarity. The higher the threshold, the more semantically similar is the tweet to one of the prototypical tweets from Table I. In comparison, in Table III, it can be observed that the seven control group tweets, when $\alpha = 0.50$, the % relative control group tweets remains low.

TABLE I. PROTOTYPICAL EVENT TWEETS AND CONTROL GROUP TWEETS

Tweet ID	Prototypical Event Tweets
1	There are young people rioting, smashing cars and vandalizing buildings.
2	The Bullring Shopping Centre has been closed amidst fears of looting and rioting. Large police presence
3	Don't understand how people think they can just hear of rioting and go down and loot! Man on news says it was older people!
4	I'm glad I'm in a peaceful country where people respect each other while the UK burns! Philippines
5	Sending in army may clear streets but it would be a sign of major political weakness for Cameron London riots
6	Don't call them anarchists. Anarchy is a political philosophy. This is just shopping with no rules. Call them capitalist
7	Rioting & looting has spread across UK – London
	Control Event Tweets
8	Welcome back Egypt Jan 25
9	Helicopter hovering above Abbottabad at 1AM (is a rare event)
10	my daughter her name is sarah m. rivera
11	This lockout is really boring..anybody playing flag football in Ok..I need to run around or something!
12	Brooms up London!
13	Here's another Photo of the shuttle from my plane.
14	Earthquake

TABLE II. %RELATIVE PROTOTYPICAL TWEETS (1 TO 7) USING FAST

Day	$\alpha = 0.50$	$\alpha = 0.55$	$\alpha = 0.60$	$\alpha = 0.65$	$\alpha = 0.70$
1	13.21	5.21	1.21	0.21	0
2	12.21	4.57	1.35	0.14	0
3	14.78	6.57	2.42	0.35	0
4	14.85	7.21	2.21	0.78	0.07
5	11.64	4.35	1.07	0.14	0
6	18.64	9	2.71	0.57	0
7	17.71	8.42	3.71	0.85	0.14
8	18.28	10.41	4.85	1.85	0.35
9	17.28	7.92	3.07	0.57	0.21
10	23.21	11.92	4.71	1.14	0
11	32.92	18.92	9.42	3.35	0.71

TABLE III. %RELATIVE CONTROL GROUP TWEETS (8 TO 14) USING FAST

Day	$\alpha = 0.50$	$\alpha = 0.55$	$\alpha = 0.60$	$\alpha = 0.65$	$\alpha = 0.70$
1	3.92	0.92	0.28	0	0
2	2.35	0.28	0	0	0
3	5.14	1.85	0.21	0.07	0
4	4.71	2.07	0.42	0.07	0
5	3.07	1.07	0.21	0	0
6	4	1.07	0.42	0.07	0.07
7	3.28	0.92	0.28	0.14	0.07
8	3.35	0.92	0.14	0	0
9	2	0.5	0.14	0	0
10	4.92	1.35	0.28	0.07	0
11	6.57	2.07	0.5	0	0

Figures 1 and 2 visually show the effect of matching the prototypical event tweets (Table I) with varying semantic threshold over the 11 day period. The y-axis shows the % of tweets relative to each result range (i.e. similarity threshold $\alpha > 0.5$ compared to $\alpha > 0.7$) for each day which was calculated in equation 1. The x-axis shows the days 1 to 11. The graph in Figure 1 shows the event clearly peaking on the 6th August and then again showing growth up to the 11th August although this is predominantly with semantic thresholds 0.65.

The graphs in figures 1 and 2 show an increase in matched tweets for higher similarity thresholds for the days of the London riots on the 6th August, which remains relatively stable to the 9th and then sharply increases which corresponds to the triggering of further riots occurring across the rest of the UK. On examining tweets in the sample from day 11, it was observed that the tweets that matched with a higher similarity threshold were focused on clean-up operations undertaken by the general public [25]. These consisted of tweets not only from influential riot Twitter users [3], but also from those who were not. Interestingly where the similarity threshold is > 0.6 (Table II) there is an increase in similarity of tweets prior to the events; from 2.7% to 3.7% on the day before the riot and the first day of rioting. The lower α , and hence the more general the semantic match, the more evident the rise between days. The obvious dip in the %relative tweets on day 5 can only be due to the sample selection.




Fig.1. %Relative Prototypical Tweets (1 To 7) Using FAST




Fig.2. %Relative Control Tweets (8 To 14) Using FAST

In order to identify the day-to-day trends, the scaled results from day p were subtracted from the scaled results from day $p-1$, where p is the day of the month. As a result, the similarity threshold trends are shown in table IV for prototypical event tweets and Table V for control group tweets.

TABLE IV. COMPARISON OF PROTOTYPICAL EVENT TWEET SIMILARITY BETWEEN DAYS

Day	$\alpha = 0.50$	$\alpha = 0.55$	$\alpha = 0.60$	$\alpha = 0.65$	$\alpha = 0.70$
1 to 2	-1	-0.64	0.14	-0.07	0
2 to 3	2.57	2	1.07	0.21	0
3 to 4	0.07	0.64	-0.21	0.43	0.07
4 to 5	-3.21	-2.86	-1.14	-0.64	-0.07
5 to 6	7	4.65	1.64	0.43	0
6 to 7	-0.93	-0.58	1	0.28	0.14
7 to 8	0.57	1.99	1.14	1	0.21
8 to 9	-1	-2.49	-1.78	-1.28	-0.14
9 to 10	5.93	4	1.64	0.57	-0.21
10 to 11	9.71	7	4.71	2.21	0.71

TABLE V. COMPARISON OF CONTROL EVENT TWEET SIMILARITY BETWEEN DAYS

Day	$\alpha = 0.50$	$\alpha = 0.55$	$\alpha = 0.60$	$\alpha = 0.65$	$\alpha = 0.70$
1 to 2	-1.57	-0.64	-0.28	0	0
2 to 3	2.79	1.57	0.21	0.07	0
3 to 4	-0.43	0.22	0.21	0	0
4 to 5	-1.64	-1	-0.21	-0.07	0
5 to 6	0.93	0	0.21	0.07	0.07
6 to 7	-0.72	-0.15	-0.14	0.07	0
7 to 8	0.07	0	-0.14	-0.14	-0.07
8 to 9	-1.35	-0.42	0	0	0
9 to 10	2.92	0.85	0.14	0.07	0
10 to 11	1.65	0.72	0.22	-0.07	0

Plotting the results of Table IV as shown in Figure 3 shows that we can identify any day-to-day rises from previous day-to-day tweeting i.e. from the 5th to the 6th of August there was a rise of 0.21% for the similarity threshold of $\alpha > 0.6$. In the days leading up to the riot it can be seen that changing the similarity threshold does not yield any significant changes in

tweets matching the prototypical event tweets. Figure 3 shows a clear spike between the 5th and 7th August around the date of the riot and also a further increase between the 9th and 11th again corresponding to further riots that were triggered and the public clean-up operation. Figure 4 visually shows the day-to-day trends of the control group tweets where it was observed that there is no significant difference over the 11 day period.

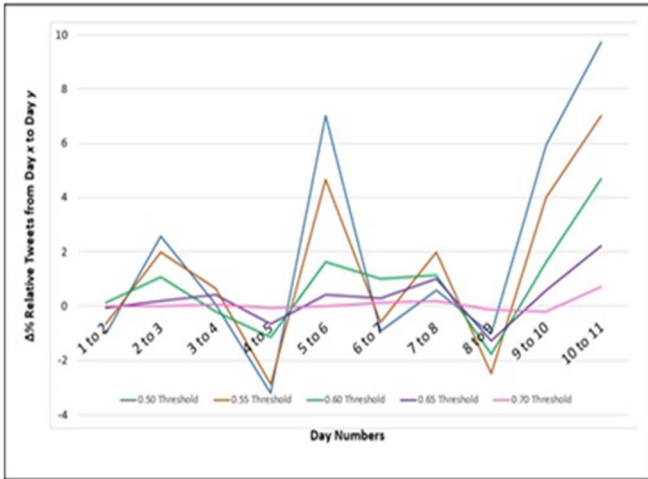


Fig 3. Identifying day-to-day trends surrounding an event (Prototypical Tweets)




Fig 4. Identifying day-to-day trends surrounding event (control tweets)

B. Discussion

The experiments have shown that an event can be potentially detected, through using a set of prototypical tweets about the event type, by measuring the rise in similarity between the prototypical event tweets. The significance is dependent on the set of prototypical tweets used and the similarity threshold of the FAST measure. It was

found that the prototypical tweets used which were extracted from the original sample of 17,795 tweets contained words that were specific to UK riots i.e. London, UK which are within Wordnet [25] which suggests that more general prototypical event tweets would need to be used to produce a more general set of event specific prototypical tweets. The results in figure 3 show that a similarity threshold of $\alpha > 0.55$ with FAST is sufficient to indicate an event, when comparing the change of percentage between two periods of time. However, a more generalised set of prototypical tweets may have yielded a high semantic similarity threshold. In comparison, there were no significant day to day trends using the control group tweets. A clear advantage of using a FSSM (identified from previous work on human correlations[17]) is that words such as ‘young’, ‘older’, ‘rare’ and ‘major’ highlighted only from the prototypical event tweets would allow their semantic meaning in the context of the tweet syntax to contribute towards the overall similarity of two comparison tweets. Therefore, produce a measurement more in line with human interpretation.

VI. CONCLUSIONS

The overall conclusion from the experiments conducted is that using a fuzzy semantic similarity measure such as FAST, makes it possible to detect potential events using fuzzy short text semantic similarity measures and prototypical event tweets. This confirms the hypothesis H1 is true. The changes in semantic similarity measurements between dates are able to indicate a potential event. This is based on the assumption that trends in rises of matches with a set of prototypical event tweets is indicative of a potential future event. However, the semantic threshold required to show a potential event was lower than expected, typically = 0.55. FAST was designed to be used on short texts such as sentences which have an established structure and when calculating the total similarity, inherited the weightings between the semantic part and the syntactic component from STASIS [23]. Given that these weights were designed through empirical experiments on structured sentences and not unstructured texts such as tweets, further work will example the weightings between the semantic, syntactic and fuzzy components within FAST. In order to validate the methodology for detecting possible events using FSSM, further experiments would need to be carried out to see if similar patterns occur with other historical events.

Many police departments around the world currently monitor "social media risk" with different degrees of success [27]. A recent report suggested that using such tools, they had the ... "potential to remove any bias from the picture presented..." [28]. the report also highlighted that "... police force representatives thought it may not be as flexible as their more qualitative approach..." and they could be improved by adding further information. Hence, further work should seek to integrate analysis of the semantic meaning of tweets into larger social network analysis systems.

REFERENCES

- [1] The statistics Portal. [online], Available at <https://www.statista.com/statistics/282087/number-of-monthly-active-Twitter-users/> [Accessed 12/1/2017], 2017.
- [2] Ifrim, G., Shi, B. & Brigadir, I. (2014), Event detection in Twitter using aggressive ltering and hierarchical tweet clustering, in 'Second ACM Workshop on Social News on the Web (SNOW), Seoul, Korea, 2014.
- [3] Morozov, E., Guardian.co.uk Facebook and Twitter are just places revolutionaries go. [Online] Available at: <http://www.guardian.co.uk/commentisfree/2011/mar/07/facebook-Twitter-revolutionaries-cyber-utopians> [Accessed 12/1/2017], 2011.
- [4] Munroe, R. [Online] Available at <http://blog.xkcd.com/2011/08/24/earthquakes/> [Accessed 11/1/2017]
- [5] Ford, R. Hollywood Reporter. [Online] Available at: <http://www.hollywoodreporter.com/news/earthquake-Twitter-users-learned-tremors-226481> [Accessed 12/1/2017]
- [6] Gupta, M. Li, R. Chang, K. Towards a Social Media Analytics Platform: Event Detection and Description for Twitter – a Tutorial, 23rd International WWW Conference, [Online] Available at: <http://www.2014.kr/asset/slide/Towards%20a%20Social%20Media%20Analytics%20Platform.pdf>, [Accessed 11/ 1/ 2017], 2014
- [7] Katragadda, S, Virani, S, Benton, R. Raghavan, V. Detection of event onset using, IEEE IJCNN, DOI: : 10.1109/IJCNN.2016.7727381, 2016.
- [8] Sakaki, Takeshi, Makoto Okazaki, and Yutaka 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. ACM, 2010.
- [9] Pavlyshenko, B. Forecasting of Events by Tweet Data Mining." arXiv preprint arXiv: 1310.3499 [Online] Available at: <http://arxiv.org/abs/1310.3499>, [Accessed 12/1/2017], 2013.
- [10] Arias , M. Arratia, A. Xuriguera , R, Forecasting with Twitter data, ACM Transactions on Intelligent Systems and Technology (TIST) - Special Section on Intelligent Mobile Knowledge Discovery and Management Systems and Special Issue on Social Web Mining archive , Vol5:1, DOI: 10.1145/2542182.2542190. 2013
- [11] Rui, L. Lei, K. Khadiwala, R. Chang, K. TEDAS: A Twitter-based Event Detection and Analysis System, IEEE 28th International Conference on Data Engineering, pp,1273 – 1276, 2012.
- [12] Silvio S. Ribeiro, Jr., Clodoveu A. Davis, Jr., Diogo Rennó R. Oliveira, Wagner Meira, Jr., Tatiana S. Gonçalves, and Gisele L. Pappa, Traffic observatory: a system to detect and locate traffic events and conditions using Twitter. In Proceedings of the 5th ACM SIGSPATIAL International Workshop on Location-Based Social Networks (LBSN '12). ACM, pp.5-11. DOI=10.1145/2442796.2442800, 2012.
- [13] Pohl, D. Bouchachia, A. Information Propagation in Social Networks During Crises: A Structural Framework, Propagation Phenomena in Real World Networks, Vol:85, Intelligent Systems Reference Library, pp 293-309, 2015
- [14] Rubenstein, H. Goodenough, j. Contextual Correlates of Synonymy. Communications of the ACM, 8, pp.627-633, 1965.
- [15] Chandran, D. Crockett, K. Bandar, Z. Mclean, D. FAST: A Fuzzy Semantic Sentence Similarity Measure, accepted for the IEEE International Conference on Fuzzy Systems, India, Digital Object Identifier :10.1109/FUZZ-IEEE.2013.6622344, 2013
- [16] O'shea, James, Zuhair Bandar, and Keeley Crockett. "A new benchmark dataset with production methodology for short text semantic similarity algorithms." ACM Transactions on Speech and Language Processing (TSLP) 10.4 (2013): 19.
- [17] Crockett, K. Chandran, D, Mclean, D. On the Creation of a Fuzzy Dataset for the Evaluation of Fuzzy Semantic Similarity Measures, IEEE WCCI – FUZZY Systems, China, pp. 752..759, DOI: 10.1109/FUZZ-IEEE.2014.6891571, 2014.
- [18] BBC, As it happened: England riots day five. [Online] Available at: <http://www.bbc.co.uk/news/uk-14449675> [Accessed 13/6/2014, 2011
- [19] BBC, 2011. Riots in Tottenham after Mark Duggan shooting protest. [Online] Available at: <http://www.bbc.co.uk/news/uk-england-london-14434318> [Accessed 13 Jun 2014], 2011.
- [20] Evans, L., 2011. 200 most influential Twitter users during the riots: are you on the list. [Online] Available at: <http://www.guardian.co.uk/news/datablog/2011/dec/08/riot-Twitter-top-200> [Accessed 13 Jun 2014], 2011.
- [21] The Guardian and LSE. Reading The Riots: Investigating England's Summer Of Disorder [Online] Available at: <http://www.guardian.co.uk/uk/series/reading-the-riots>. [Accessed 13th June 2014], 2011.
- [22] Tonkin, E. Pfeiffer, H. Tourte, G. Twitter, Information Sharing and the London Riots? by [Online] Available: https://www.asis.org/Bulletin/Dec-11/DecJan12_Tonkin_Pfeiffer_Tourte.pdf. [Accessed 12/1/2017], 2014.
- [23] Li, Y. Mclean, D. Bandar, Z. O'Shea, J. Crockett, K. "Sentence similarity based on semantic nets and corpus statistics", IEEE Transactions on Knowledge and Data Engineering, vol. 18, no. 8, pp.1138-1150, 2006.
- [24] Chandran, D, Crockett, K. Mclean, D, An Automatic Corpus Based Method for a Building Multiple Fuzzy Word Dataset, IEEE-FUZZ 2015, DOI: 10.1109/FUZZ-IEEE.2015.7337877, 2015.
- [25] Princeton University, About Wordnet. [Online] Available at: <http://wordnet.princeton.edu/> [Accessed 13 June 2014], 2017.
- [26] ABC news. The Year According to Twitter: 2011's Top Tweets Available [Online] Available: <http://abcnews.go.com/Technology/year-Twitter-2011s-top-tweets/story?id=15065335> [Accessed 12/1/2017], 2012.
- [27] As Police Monitor Social Media, Legal Lines Become Blurred, by Martin Kaste, February 28, 2014 8:39 PM ET, [Online] Available: <http://www.npr.org/blogs/alltechconsidered/2014/02/28/284131881/as-police-monitor-social-media-legal-lines-become-blurred> [Accessed 22 January 2017], 2014.
- [28] Gunnell, D. Hillier, J. Blakeborough, L. Social Network Analysis of an Urban Street Gang Using Police Intelligence Data, Research Report 89, The Home Office. [Online]. Available at: https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/491578/horr89.pdf. [Accessed 22/01/2016], 2016.