

Extraction of unexpected rules from Twitter hashtags and its application to sport events

Conference or Workshop Item

Accepted Version

Adedoyin-Olowe, M., Gaber, M. M., Dancausa, C. M. and Stahl, F. (2014) Extraction of unexpected rules from Twitter hashtags and its application to sport events. In: 13th International Conference on Machine Learning and Applications (ICMLA 2014), 3-5 Dec 2014, Detroit, MI, USA, pp. 207-212. Available at <http://centaur.reading.ac.uk/40751/>

It is advisable to refer to the publisher's version if you intend to cite from the work. See [Guidance on citing](#).

Published version at: <http://dx.doi.org/10.1109/ICMLA.2014.38>

All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in

the [End User Agreement](#).

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

Extraction of Unexpected Rules from Twitter Hashtags and its Application to Sport Events

Mariam Adedoyin-Olowe, Mohamed Medhat Gaber
and Carlos Martin Dancausa
School of Computing Science and
Digital Media

Robert Gordon University, Aberdeen, United Kingdom
AB10 7GJ, UK

Email:m.a.adedoyin-olowe@rgu.ac.uk
m.gaber1@rgu.ac.uk
c.j.martin-dancausa@rgu.ac.uk

Frederic Stahl

School of Systems Engineering, University of Reading
PO Box 225, Whiteknights, Reading, RG6 6AY, UK
Email:F.T.Stahl@reading.ac.uk

Abstract—Twitter has become a dependable microblogging tool for real time information dissemination and newsworthy events broadcast. Its users sometimes break news on the network faster than traditional newsgroups due to their presence at ongoing real life events at most times. Different topic detection methods are currently used to match Twitter posts to real life news of mainstream media. In this paper, we analyse tweets relating to the English FA Cup finals 2012 by applying our novel method named *TRCM* to extract association rules present in hashtag keywords of tweets in different time-slots. Our system identify evolving hashtag keywords with strong association rules in each time-slot. We then map the identified hashtag keywords to event highlights of the game as reported in the ground truth of the main stream media. The performance effectiveness measure of our experiments show that our method perform well as a Topic Detection and Tracking approach.

I. INTRODUCTION

News and event-centred data sharing and searching on Twitter has become very common in recent times. Twitter subscribers and other entities such as events organisers, businesses and government organisations encourage their audience to post tweets about their opinions regarding specific event by creating official hashtags for the event to keep their audience informed. Events and other occurrences attract attention on Twitter which result in these events becoming popular, while also attracting attention from different quarters. Private individuals and other entities are now relying more on information posted on the network for decision-making. Twitter affords its users the opportunity of sharing and receiving news/information in real time. Traditional newsgroups follow the activities on Twitter to assist them in updating related news [15]. Many Topic Detection and Tracking *TDT* researches have been conducted on Twitter data [1], [2], [6], [7], [13], [14], [16]–[18], [22], [25], the results of which show the dynamism of Twitter network as information dissemination tool and the efficiency of data mining techniques in extracting newsworthy contents from tweets posted online.

Twitter is known to have played major role in publicising and sustaining the trend of major news of global relevance on the network. Events like the 2011 Egyptian political uprisings [21]

and the US Elections 2012 [24] were widely tweeted giving Twitter users not only the opportunity to air their views, but to also follow the trend of these events. Tweets’s hashtag (#) is a prefix symbol attached to prominent keywords in tweets. Hashtags define tweets’ content and enhances tweets’ readability. According to Laniado & Mika [11], hashtags specification is the most effective way to extract tweets of a particular topic on Twitter.

In this paper we employ our novel method named *TRCM* (Transactional-based Rule Change Mining) to detect landmark events in the English FA Cup Finals 2012 at (near) real time. We also compute the burst in the volume of tweets posted by the teams’ fans shortly after each goal.

We extend our experiments in [1], [9] by applying *TRCM* on a dynamic sport datasets; the English FA Cup Finals 2012. Our method first detects frequent hashtags present in the related tweets, which we then apply as parameter for *TDT* as contained in the ground truth source utilized. As far as we know, Association Rule Mining (*ARM*) has not been used on tweets hashtags for *TDT*. *TRCM* was able to detect landmark events that occurred during the 90 minutes (plus 6 minutes extra time) of the game played on the May 5, 2012 between Chelsea football club and Liverpool football club. English FA cup (<http://www.thefa.com/thefacup>) is a popular and widely viewed football tournament within and outside Europe.

The rest of the paper is organised as follows: Section II discusses other well performed topic detection methods already employed on Twitter data. Section III presents a brief description of *ARM* and definitions of *TRCM* rules as well as our Rule Matching concept. Section IV explains our methodology, while Section V demonstrates our experiment and we present our experimental result in Section VI. We conclude the paper in section VII with a discussion and future work.

II. RELATED WORK

TDT research is dated back to 1997 [4]. *TDT* methods extract interesting topics from Twitter streaming data and present patterns that demonstrate a representation of specific real life topics. This is often achieved by mapping detected

results to real life news/events and subsequently tracking the evolvments of such topics. Since Twitter streams high volume of data very rapidly, it is important to apply TDT on Twitter data in order to extract relevant topics from the network.

Incremental online clustering and filtering framework can be used to distinguish between messages about real life events and non-events [6]. This framework clusters subsequent tweet-based message similarity with existing clusters. On the other hand, graph-based approaches can detect keyword clusters in tweets based on their pairwise comparison. This can be a term unison graph with nodes clustered and the use of community detection algorithm based on betweenness centrality [19]. The authors of [10] used a clustering technique to detect events using a text classifier. Swit Phuvipadawat Murtata [17] proposed a method of collecting, grouping, ranking and tracking breaking news in Twitter. They built a framework named 'Hoetstream' to enable users discover breaking news from Twitter timeline. Other approaches considered first story detection on the network, [16] presented a method that merges Twitter and Wikipedia in order to enhance event detection. They explore the latency between the two streams and discovered that Twitter is more up-to-date in real life events posting. Corney et al [8] used n-grams and $df - idf_t$ to identify the most bursty terms and then using hierarchical clustering, they identified term clusters whose similarities are high as a representation of the same topic. They merged clusters to the point where each cluster is assumed to signify a distinct topic. They presented a more comprehensive detail of their algorithm in [3] by identifying real world topic in the 2012 US Presidential Elections, the US Super Tuesday primaries in 2012 and the English FA Cup 2012.

Unlike approaches reviewed in this paper, our method is able to detect unexpected events that occur in quick successions (as in football events) by identifying Association Rules ARs present in tweets' hashtags. Where other methods distinguish events tweets from non-event tweets by clustering [6], [10], [19], our method detects and lists related hashtags of event highlights as they unfold. However, our experiments in this paper are similar to Corney et al [8] but our methods are different. We apply our novel method $TRCM$ and $apriori$ of ARM to discover *unexpected rules* in hashtag keywords in tweets posted during the FA Cup Finals 2012 to detect event highlights in the match. The event highlights in real life are represented by hashtag keywords detected by our system as unexpected rules. Note that AR is in the form $X \Rightarrow Y$, where X and Y are disjoint sets of items. The overview of ARM and definition of $TRCM$ rules is presented in Section III.

III. OVERVIEW OF ASSOCIATION RULE MINING

ARM is a data mining technique used for extracting remarkable associations between different arrays of items in transactional databases (Market Basket), relational databases (Personal Details), or any other information warehouses [12] in the form of rules. It analyses the frequent if/then (antecedent/consequent) patterns by using the support and confidence measures to discover most significant relationships [5]. For instance, if a buyer purchase bread, there is 90% chances that he will also purchase butter. *Support* indicates how frequently the items appear in the database while *confidence* indicates the number of times the if/then statements is positive. *Apriori* algorithm of ARM discovers frequent itemsets

and strong association rules. This technique is capable of unveiling every feasible association that satisfies the user-defined threshold of support and confidence.

A. Rule Dynamics of Association Rule Mining

$TRCM$ (Transaction-based Rule Change Mining) is a methodology that applies ARM to extract evolving ARs in tweets' hashtags at two consecutive periods. The detected ARs are subsequently mapped to specific evolving news/events in real life. The aim of $TRCM$ is to demonstrate the relevance of hashtags in online tweets and the importance of appropriate use of hashtags in tweets.

In our previous paper [1] we employed the left hand side *lhs/conditional* and the right hand side *rhs/consequent* parts of rules in *Apriori* principle to analyse hashtags present in tweets. The analysis of the *lhs* and the *rhs* is then used to detect the ARs present in tweets at different points in time. The similarities and differences in the ARs in the rulesets r_i^t and r_j^{t+1} (where t is the time and i/j are rules present in tweets at t and $t+1$ respectively) are measured to disclose the different rules namely 'emerging', 'unexpected', 'new', and 'dead' rules in tweets Fig. 1.




Fig. 1: $TRCM$ Process

B. Notation

We employ the method of Liu et al [12] and Song et al [20] for calculating similarities and differences between two rules at 2 different period. Given that the methods implemented in [12] and [20] were aimed for ARs detected from relational datasets (general association rules). Our method defines the similarity based on the principles of degree of similarity proposed in [12] and [20]. Details of the calculations and notation used are stated as follows:

n	Number of hashtags
i	an association rule in r_i^t presented in binary vector
j	an association rule in r_j^{t+1} presented in binary vector
lh_i/lh_j	number of hashtags with value 1 in conditional part of rule i/j
rh_i/rh_j	number of hashtags with value 1 in consequent part of rule i/j
lh_{ij}/rh_{ij}	number of same hashtags in conditional/consequent part of rules i and j
p_{ij}/q_{ij}	degree of similarity of hashtags in conditional/consequent part of rules i and j
r_i^t	Rule i present at time t
r_j^{t+1}	Rule j present at time $t+1$

C. Measuring Similarity

$$p_{ij} = \frac{lh_{ij}}{\max(lh_i, lh_j)} \quad (1)$$

$$q_{ij} = \frac{rh_{ij}}{\max(rh_i, rh_j)} \quad (2)$$

Equations 1 and 2 were adopted from Liu et al [12]. They demonstrate the degree of similarity in the conditional p_{ij} and consequent q_{ij} parts of rule i and rule j at different time t and $t + 1$ in that order. *TRCM* rules are identified by matching all rules in the left hand side *lhs* and right hand side *rhs* of ruleset $t + 1$ with those in *lhs* and the *rhs* of ruleset t using the similarity and difference comparison threshold thp_{ij} and thq_{ij} . Change in rules can be discovered by matching every rule in ruleset r_i^t with those in ruleset r_j^{t+1} using the similarity and difference comparison as presented in Fig. 2 and explained in the following algorithm.

Step 1: For each rule in the conditional part of ruleset r_j^{t+1} (new), match with all rules in the conditional part of ruleset r_i^t (old). Compute the number of hashtags with value 1 (similar) present in the conditional parts of r_j^{t+1} and r_i^t .

Step 2: Divide the computed figure in step 1 by the maximum number of hashtags with value 1 in the conditional parts of either r_j^{t+1} and r_i^t .

Step 4: Identify the degree of similarity of rules in the old and new rules. However, for two rules to be similar, their degree of similarity must be greater than the pre-defined RMT (0).




Fig. 2: TRCM Rules Assignment

D. Rule Matching

In our previous experiments [1] and [9] we used Rule Matching (*RM*) to detect change in patterns of rules in r_i^t

and r_j^{t+1} , where t is the time and i/j are the rules present in the tweets. Rules in r_j^{t+1} are matched against rules in r_i^t to detect patterns of rule changes. We set the *RM* threshold in our experiments in the interval of binary vector $[0, 1]$, where 0 specifies difference and 1 specifies similarity.

E. TRCM Rules Definition

Matching rules similarities in r_i^t and r_j^{t+1} allow us to define *TRCM* rule change patterns. **Unexpected Consequent rule** occurs where a rule in r_i^t and r_j^{t+1} has similar conditional part but different consequent part ($p_{ij} \geq thp_{ij}$ and $q_{ij} < thq_{ij}$). **Unexpected Conditional rule** is detected when the consequent parts of rule r_i^t and r_j^{t+1} are similar, but the conditional parts are different ($p_{ij} < thp_{ij}$ and $q_{ij} \geq thq_{ij}$). **Emerging rules** are discovered when rules at time t and $t + 1$ have similarities in both the conditional and consequent parts of the rule that are greater than the user-defined threshold ($p_{ij} \geq thp_{ij}$ and $q_{ij} > thq_{ij}$). An example of an emerging rule occurrence in real life scenario is the breaking news of a disaster like earthquake in Japan [1]. Tweets relating to breaking news increase in volume very rapidly resulting in emerging rules within a short period. However, all rules are said to be **new** until there is a matching rule found ($p_{ij} < thp_{ij}$ and $q_{ij} < thq_{ij}$). Every rule at time $t + 1$ is completely different from all the rules in time t . New rules are found to be frequent but have no matching rule in t . A **‘Dead’ Rules** occurrence is the opposite of new rules detection. A rule in t is labeled ‘dead’ if its maximum similarity measure with all the rules in $t + 1$ is less than the user-defined threshold ($p_{ij} < thp_{ij}$ and $q_{ij} < thq_{ij}$). ‘Dead’ rules in real life are topics that were initially present on Twitter but disappear after some time. In our experiments we consider only the unexpected rules as they represent the occurrence of event highlights during the football match. The hashtag keywords in the unexpected rules (consequent and conditional) are then mapped to the event highlights presented in the ground truth to validate our experiments.

IV. METHODOLOGY

The game featured Chelsea Football Club and Liverpool Football Club, with both teams having huge amount of fans in and outside the UK. Expectedly, fans of these two football clubs went on Twitter to report events of the match as they unfold. We train our system to discover unexpected rules in tweets hashtags of the FA Cup game, which we then apply to the real life event highlights of the game as contained in the ground truth. We divide the entire tweets retrieved from Twitter API into window size and select an update period of 1 minute as in Corney et al [8]. We selected 1 minute update time considering the dynamic nature of football matches where events evolve very rapidly. This setting enhances the precision of rules returned by *TRCM* within each time-slot. We also set both the support and confidence to 0.001 after carrying out preliminary study to confirm the setting that best optimise *TRCM* efficiency on the case study dataset. We choose a low minimum support to eliminate the problem with using global minimum support threshold that often misses the not-so-frequent but important items in the datasets. We extract all hashtags in tweets and define a function that rapidly finds

equal terms in *lhs* and *rhs* (left-hand side and right-hand side) rules of r_i^t and at r_j^{t+1} . These are used to set the *TRCM* rules. We find matching values in *lhs* and *rhs* of r_j^{t+1} and r_i^t as presented in Fig. 3 where *#tcot* and *#RonPaul* are unexpected consequent rules (similar lhs but different rhs) and *#Romney* is a new rule (no matching found). We then convert the value to similarity by dividing the value of *lhs/rhs* by the total number of *lhs/rhs* in r_i^t . *TRCM* is identified by defining the thp_{ij} and thq_{ij} (left hand side and right hand side user-defined threshold) which are set between 0 and 1.

lhs	rhs	support	confidence	lift
1 {#RonPaul}	=> {#TeaParty}	0.0235	0.4795918	17.43970
2 {#TeaParty}	=> {#tcot}	0.0180	0.6545455	10.14799
3 {#tcot}	=> {#TeaParty}	0.0180	0.2790698	10.14799
4 {#gop2012}	=> {#RonPaul}	0.0205	0.6833333	13.94558
5 {#RonPaul}	=> {#gop2012}	0.0205	0.4183673	13.94558
6 {#gop2012}	=> {#tcot}	0.0270	0.9000000	13.95349

Rulesets at r_i^t

lhs	rhs	support	confidence	lift
1 {#gop2012}	=> {#tcot}	0.0270	0.9000000	13.953488
2 {#tcot}	=> {#gop2012}	0.0270	0.4186047	13.953488
3 {#Romney}	=> {#Santorum}	0.0205	0.5256410	11.066127
4 {#Santorum}	=> {#Romney}	0.0205	0.4315789	11.066127
5 {#RonPaul}	=> {#tcot}	0.0215	0.4387755	6.802721
6 {#tcot}	=> {#RonPaul}	0.0215	0.3333333	6.802721
7 {#newt2012}	=> {#AK}	0.0170	0.9444444	53.968254

Rulesets at r_i^{t+1}

Fig. 3: Rules Matching Sample

The total *TRCM* extracted from the collection of tweets analysed is evaluated by dividing the total number of tweet messages (*TM*) by the time-slot (*TS*) window-size.

$$\frac{TM}{TS} = TRCM$$

The events mapped by our system include goals, bookings, substitutions, shot-on-targets, free-kicks and foul plays. For the experiments, an item *h* is any hashtag present in the tweet, while the transaction is the tweet message that occurs in a time slot *T*. The number of times that any given set of hashtags occurs in the time slot is referred to as its *support*, and itemset (hashtag) that meets a minimum support is refer to as a *pattern*.

V. EXPERIMENT

A. Data Collection and Preprocessing

We crawl the Twitter streaming API and collected all tweets on FA cup final 2012 [8] using the official hashtags for the event (*#FACupFinal2012*), the two team names (*#CFC*, *#LFC*, *#ChelseaFC*, *#LiverpoolFC*) and names of key players. We normalize all the tweets by removing their meta-data such as user ID, user mentions, location and url pointers. We retain tweets' timestamps as shown in Fig. 4, the timestamps enable us to map hashtag keywords detected by *TRCM* to events in the ground truth in (near) real time. Although we collected 444,291 tweets over a period of 72 hours (4 - 6 May 2012), we analyse only the 224,291 tweets that were posted within the duration of the match (May 5 2012, 5:15pm to 7:00pm).

Since hashtags usage on Twitter network is to describe tweets' content, we set out to carry out the *TDT* experiments

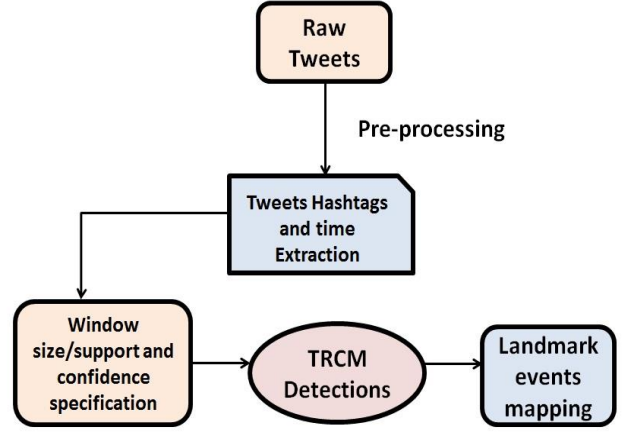


Fig. 4: Event Detection Process

in this paper to identify the efficiency of *ARM* technique on hashtag keywords in dynamic football event tweets. We achieve this by extracting *ARs* in hashtag keywords in Twitter dataset of *FA Cup Finals 2012* at different time slots during the course of the match. To detect events highlights occurrence during the football match we examine the detected hashtags keywords in each time-slot in the entire stream of tweets (224,291) and then rank changes in *ARs* to find the *unexpected consequent* and *unexpected conditional rules* in each time-slot. Hashtags detected in each time-slot are recorded along with the time the tweet was posted on Twitter to evaluate our system. We establish a detection if the returned hashtags in each time-slot contain at least one out of key terms used in the ground truth within the same time frame the detection occurred.

We confirm that the hashtag keywords detected as unexpected rules are those that best represent different event highlights during the football match when mapped to the ground truth as shown in Table I. The annotation is done manually by mapping the hashtags detected within specific time-slot to events in the ground truth. During the annotation, we evaluate the number of tweets posted by fans of the two teams in the time-slot of each goal by counting the number of times *#CFC* and *#LFC* occur within these time-slots. The statistic shows that *CFC* dominated Twitter network after each goal including the third goal scored by Liverpool Club player in the 64th minute of the match. We measure *TRCM* performance against the method used in the benchmark. Detailed comparison is discussed in section VI.

VI. EXPERIMENTAL RESULT

To validate our topic detection technique we generate a ground truth from the BBC sport live text commentary from their official website (<http://goo.gl/Ir3Of>). The FA Cup final match between *Chelsea Football Club* and *Liverpool Football Club* produce event highlights that were detected by *TRCM*. Our system was able to detect events such as **goals scored, bookings, player substitutions, free kicks, offside, misses, saves and clearances**. Events detection mapping was carried out manually as presented in Table I. A graph showing the performance profile of *TRCM* against the ground truth is

presented in Fig. 5. We also show the statistic of fans' tweets burst after each goal in Fig. 6. This was obtained by manually counting the number of times the names of each of the clubs were hashtagged within the time window in which each goal were scored. We present a graph showing *TRCM* performance against the benchmark in Fig. 7. The effectiveness of measure for our system is discussed in the next section.

TABLE I: Table Showing TRCM Event Detection Mapping

	G	Sub	BK	FK	S	CL	OS	BL	MS	TE
GT	3	4	3	10	11	19	4	3	2	59
TRCM	3	3	1	10	9	19	4	3	2	54

In Table.1 GT = Ground Truth; G = Goals; Sub = Substitutions; BK = Bookings; FK = Free kicks; CL = Clearances; OS = Offsides; BL = Blocks; MS = Misses; TE = Total Event




Fig. 5: TRCM /Ground Truth Performance Profile




Fig. 6: Statistic of Tweets Burst After Goals

A. Effectiveness Measure

Recall and precision are commonly used in Information Retrieval (*IR*) experiments to evaluate the effectiveness of the system. Precision is the percentage of relevant instances in the datasets returned by the system. On the other hand, recall is the percentage of relevant instances (hashtags) classified correctly. In *IR*, the system error rates are used to evaluate significance of the system. Numerous approaches of single-valued measures

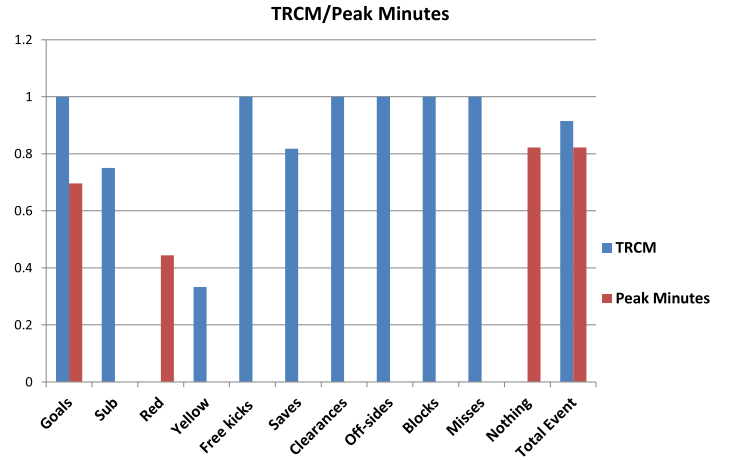


Fig. 7: TRCM Performance Graph

have been implemented as reviewed by van Oorschot et al [23]. However, *F - Measure*, a combination of precision and recall, stands out as a dominant approach in evaluating text classification.

In this experiment, we made use of the *F - Measure*; where *B* indicates that the system produces a false alarm event, *C* indicates that the system produces a miss and *A* indicates that a hashtag is relevant to an event. We summarize the method by calculating the effectiveness measures of our system as shown in Table II.

TABLE II: Table showing the effectiveness measure of TRCM

	Retrieved	Relevant	Not-relevant
	A = 54	B = 9	-
Not retrieved	C = 5	-	-

The retrieved hashtags in the table are those that our system classified as positive instances to the football game, and the relevant hashtags are those that we manually judged relevant to the game. The effectiveness measure of our system implies that at least one out of the hashtags detected must be included in the relevant slot of the ground truth minute-by-minute commentary.

$$Precision = P = \frac{A}{A + C} \times 100 = 91.5\%$$

$$Recall = R = \frac{A}{A + B} \times 100 = 85.7\%$$

$$F-Measure = F = \frac{2PR}{P + R} = 88.5\%$$

VII. DISCUSSION AND FUTURE WORK

We use the approach of van Oorschot et al [23] as a benchmark for our experiments for the following reasons 1) both games are similar in nature, van Oorschot et al [23] collected Dutch premier league soccer related tweets while we collected English FA Cup Final tweets. 2) both experiment

applied hashtags in detecting events in the games. 3) we both work on similar volume of datasets - In the experiment in [23] a total of 315,844 tweets were analysed, while in our experiments we analysed a total of 224,291 tweets. However, our method is different, they chose three sets of game minutes namely; ALL MINUTES, where no filtering is performed (used as baseline), PEAK MINUTE, where only minutes from the best performing peak selection method are included and EVENT MINUTES are only minutes in which an event takes place. Their system was developed to detect events that occur in these selected minutes. They consider five classes of events including: goals, own-goals, red cards, yellow cards and substitutions. In our experiments, we consider additional events like; (goal) misses, (goal) clearances, (goal) saves and free-kicks, which are believe to form an integral part of any football/soccer match. They used 10-fold cross-validation of Weka Toolkit and the LibSVM library where SVM performed best with the F-measure results. Our experimental analysis also utilised F-score (otherwise known as F-measure). We measure our system against their Peak Minute Only as this selection is meant to return event occurrences at peak minutes. One of the down side of their selection is the window size which they set at 5 minutes interval, with this setting, 29% of their selected peak minutes fail to detect any of the 5 classes of events they set out to detect. From our experiment we noticed that the more window size is increased the lesser the chances of our system detecting event highlights in the match. It is evident in Fig. 7 that *TRCM* event detection precision out-performed LibSVM employed in [23] by **0.915** to **0.822**. In addition, *TRCM* returned more event highlights than the method used in [23]. While reviewing related work already carried out in the area of *TDT* in Twitter, we notice that hashtags as Twitter object and *ARM* as data mining techniques has not gained popularity in *TDT* on Twitter even though tweets' hashtags are meant to describe tweets' contents and enhance the readability of tweets. Our experimental results shows that tweet hashtags (if appropriately used) are capable of conveying events highlights in sport games in real time. In our future experiments, we shall be analysing datasets from different domains such as politics and entertainment in order to compare *TRCM* performance on datasets from diverse domains and investigate how aggregate features of datasets from these domain affect *TRCM* performance.

REFERENCES

- [1] M. Adedoyin-Olowe, M. M. Gaber, and F. Stahl. Trcm: a methodology for temporal analysis of evolving concepts in twitter. In *Artificial Intelligence and Soft Computing*, pages 135–145. Springer, 2013.
- [2] P. Agarwal, R. Vaithyanathan, S. Sharma, and G. Shroff. Catching the long-tail: Extracting local news events from twitter. In *ICWSM*, 2012.
- [3] L. M. Aiello, G. Petkos, C. Martin, D. Corney, S. Papadopoulos, R. Skraba, A. Goker, I. Kompatsiaris, and A. Jaimes. Sensing trending topics in twitter. 2013.
- [4] J. Allan. Introduction to topic detection and tracking. In *Topic detection and tracking*, pages 1–16. Springer, 2002.
- [5] Anonymous. Association rules (in data mining).searchbusiness analytics, retrieved from <http://goo.gl/avvnfj> on 23.6.14.
- [6] H. Becker, D. Iter, M. Naaman, and L. Gravano. Identifying content for planned events across social media sites. In *Proceedings of the fifth ACM international conference on Web search and data mining*, pages 533–542. ACM, 2012.
- [7] H. Becker, M. Naaman, and L. Gravano. Beyond trending topics: Real-world event identification on twitter. *ICWSM*, 11:438–441, 2011.
- [8] D. Corney, C. Martin, and A. Göker. Spot the ball: Detecting sports events on twitter. In *Advances in Information Retrieval*, pages 449–454. Springer, 2014.
- [9] J. B. Gomes, M. Adedoyin-Olowe, M. M. Gaber, and F. Stahl. Rule type identification using trcm for trend analysis in twitter. In *Research and Development in Intelligent Systems XXX*, pages 273–278. Springer, 2013.
- [10] A. Jackoway, H. Samet, and J. Sankaranarayanan. Identification of live news events using twitter. In *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on Location-Based Social Networks*, pages 25–32. ACM, 2011.
- [11] D. Laniado and P. Mika. Making sense of twitter. In *The Semantic Web–ISWC 2010*, pages 470–485. Springer, 2010.
- [12] D.-R. Liu, M.-J. Shih, C.-J. Liau, and C.-H. Lai. Mining the change of event trends for decision support in environmental scanning. *Expert Systems with Applications*, 36(2):972–984, 2009.
- [13] M. Mathioudakis and N. Koudas. Twittermonitor: trend detection over the twitter stream. In *Proceedings of the 2010 ACM SIGMOD International Conference on Management of data*, pages 1155–1158. ACM, 2010.
- [14] M. Naaman, H. Becker, and L. Gravano. Hip and trendy: Characterizing emerging trends on twitter. *Journal of the American Society for Information Science and Technology*, 62(5):902–918, 2011.
- [15] N. Newman. The rise of social media and its impact on mainstream journalism. *Reuters Institute for the Study of Journalism*, 2009.
- [16] M. Osborne, S. Petrovic, R. McCreadie, C. Macdonald, and I. Ounis. Bieber no more: First story detection using twitter and wikipedia. In *Proceedings of the Workshop on Time-aware Information Access. TAI*, volume 12, 2012.
- [17] S. Phuvipadawat and T. Murata. Breaking news detection and tracking in twitter. In *Web Intelligence and Intelligent Agent Technology (WI-IAT), 2010 IEEE/WIC/ACM International Conference on*, volume 3, pages 120–123. IEEE, 2010.
- [18] 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*, pages 851–860. ACM, 2010.
- [19] H. Sayyadi, M. Hurst, and A. Maykov. Event detection and tracking in social streams. In *ICWSM*, 2009.
- [20] H. S. Song, S. H. Kim, et al. Mining the change of customer behavior in an internet shopping mall. *Expert Systems with Applications*, 21(3):157–168, 2001.
- [21] K. Starbird and L. Palen. (how) will the revolution be retweeted?: information diffusion and the 2011 egyptian uprising. In *Proceedings of the acm 2012 conference on computer supported cooperative work*, pages 7–16. ACM, 2012.
- [22] A. Tumasjan, T. O. Sprenger, P. G. Sandner, and I. M. Welp. Predicting elections with twitter: What 140 characters reveal about political sentiment. *ICWSM*, 10:178–185, 2010.
- [23] G. van Oorschot, M. van Erp, and C. Dijkshoorn. Automatic extraction of soccer game events from twitter. In *Proc. of the Workshop on Detection, Representation, and Exploitation of Events in the Semantic Web*, 2012.
- [24] H. Wang, D. Can, A. Kazemzadeh, F. Bar, and S. Narayanan. A system for real-time twitter sentiment analysis of 2012 us presidential election cycle. In *Proceedings of the ACL 2012 System Demonstrations*, pages 115–120. Association for Computational Linguistics, 2012.
- [25] K. Watanabe, M. Ochi, M. Okabe, and R. Onai. Jasmine: a real-time local-event detection system based on geolocation information propagated to microblogs. In *Proceedings of the 20th ACM international conference on Information and knowledge management*, pages 2541–2544. ACM, 2011.