

SMERC: Social media event response clustering using textual and temporal information

Peter Mathews*, Caitlin Gray†, Lewis Mitchell‡§, Giang T. Nguyen¶, and Nigel G. Bean||

School of Mathematical Sciences,
ARC Centre of Excellence for Mathematical & Statistical Frontiers (ACEMS),
The University of Adelaide
SA 5005, Australia

*peter.mathews@adelaide.edu.au, †caitlin.gray@adelaide.edu.au,
‡Data to Decisions CRC stream lead, §lewis.mitchell@adelaide.edu.au,
¶giang.nguyen@adelaide.edu.au, ||nigel.bean@adelaide.edu.au

Abstract—Tweet clustering for event detection is a powerful modern method to automate the real-time detection of events. In this work we present a new tweet clustering approach, using a probabilistic approach to incorporate temporal information. By analysing the distribution of time gaps between tweets we show that the gaps between pairs of related tweets exhibit exponential decay, whereas the gaps between unrelated tweets are approximately uniform. Guided by this insight, we use probabilistic arguments to estimate the likelihood that a pair of tweets are related, and build an improved clustering method. Our method *Social Media Event Response Clustering (SMERC)* creates clusters of tweets based on their tendency to be related to a single event. We evaluate our method at three levels: through traditional event prediction from tweet clustering, by measuring the improvement in quality of clusters created, and also comparing the clustering precision and recall with other methods. By applying SMERC to tweets collected during a number of sporting events, we demonstrate that incorporating temporal information leads to state of the art clustering performance.

Index Terms—Twitter, clustering, social sensing, event detection, text summarization

Postings to social media platforms are increasingly being used to extract useful information about real-world events. Journalists use platforms such as Twitter to summarize breaking news stories [1], governments are interested in mining social media to provide early warning of events such as disease outbreaks [2], [3], civil unrest events [4], [5] and even natural disasters such as earthquakes [6]. Developing methods to summarize the large volumes of information generated on social media by such events is therefore of great importance for scientists and end-users alike, and an extensive amount of literature has appeared on real-time microblog summarization for event detection.

Sporting events provide a unique testing ground for event detection, due to the collective attention of often large audiences on the in-game action. While event detection techniques can be applied to any type of event, sporting events are ideal to collect and study as the time, location and hashtags are generally known beforehand. Furthermore, online reaction to sporting events can be useful case studies on human behaviour, with heavy levels of engagement and emotion generally organized around the success of opposing teams.

In this work we develop a new method for clustering mi-

croblog posts authored in response to real-world events. Distinguishing it from previous methods, our proposed methodology uses both the timing and content of messages to estimate a probability that messages are related, giving more meaningful event clusters than previous methods. Through a detailed investigation of the distributions of related and unrelated messages around particular events, we arrive at a probabilistic method for clustering messages. This approach also provides insight into the mechanisms causing a social media response to a particular event and allows us to model the distribution of response times.

Our work makes the following new contributions:

- Showing that the time gaps between pairs of related messages are exponentially distributed.
- Presenting a novel way to improve tweet clustering by incorporating both temporal and textual information.

We apply our method *Social Media Event Response Clustering (SMERC)* to three Twitter datasets collected from cricket and Australian Rules Football matches, and demonstrate that the clusters around events obtained are meaningful and that our event detection improves over existing techniques.

I. PREVIOUS WORK

There exists a broad literature on microblog post summarization for both trend and event detection, particularly for sporting events. Tweet clustering is often focused on using tweet content and user features, but has limitations based on the amount of information contained in a tweet. Studies typically use either *textual* features such as tracking the number of pre-defined keywords or hashtags, or purely *temporal* features such as the timing of posts. There exist only a small number of studies that utilise both textual and temporal features in a meaningful way [7].

Yang and Leskovec [8] summarized two key components for time series clustering in online media, a distance measure and a clustering algorithm. The most commonly used distance measure is Euclidean distance, which has been used in a variety of works, e.g. [9], [10]. More advanced measures include Dynamic Time Warping [11] and the Longest Common Subsequence [12]. The most common clustering algorithm

continues to be k-means clustering [13] despite the limitations of having to specify the number of clusters beforehand, and being sensitive to the starting point.

Leskovec *et al.* [14] analyzed how popularity of memes varies over time. Our work focuses on a similar problem, but over much shorter time scales. Memes tend to be popular for months while Twitter users tend to respond to events or tweets in seconds or minutes. This makes our context arguably more challenging.

Most apparently similar to this work, Gillani *et al.* [15] provided a way to identify key events in sporting contests by clustering both temporal and textual features. They use a threshold technique to determine whether an event is sufficiently significant, and incorporate post time as a feature in a relatively straightforward manner, by appending it to a vector of word counts collected within a window. As this method uses k-means for clustering and it is not possible to know the number of key events beforehand, it is perhaps more suited to the problem of *post-hoc* microblog summarization rather than real-time event detection. Conversely, as we show, our approach will be more suited to real-time tweet clustering and event detection; by reasoning more probabilistically about the distributions of times between tweets around events, we develop a mathematical model for incorporating temporal information into tweet clustering instead of simply adding it as part of a feature vector.

There exist many methods for using Twitter for event detection. Twevent [16] extracts continuous and non-overlapping word segments, and then calculates bursty event segments within a fixed length window. To evaluate their method, they defined ‘precision’ as the proportion of detected events related to realistic events, and ‘recall’ as the proportion of realistic events detected from the data set. Due to the difficulty in assessing whether or not a particular tweet, or cluster of tweets, is related to an event or what constitutes a ‘realistic’ event, such definitions are simultaneously important and hard to define. For comparison with previous work we employ the same definitions as used to evaluate Twevent in this study, but will also examine the quality of the clusters associated with individual events on a tweet-by-tweet level through studying the content of example tweet clusters.

Several authors have used Twitter as a method for social sensing in sporting and other events, e.g. [15], [17], [18]. Zhao *et al.* [17] used the Twitter response to events in NFL games to identify key game events such as touchdowns, interceptions, fumbles and field goals. The temporal component of their event detection method was based on post rate in time windows, using the fact that social media activity increases heavily after key events. They also used content-based keywords to determine whether keyword frequency was above a pre-defined threshold.

Tweet intensity is generally modelled as a self exciting temporal point process [19], [20]. Events and other tweets tends to drive other activity, leading to burstiness. The rate of twitter activity is affected by other factors such as human prioritization of tasks and circadian cycles [21], [22].

II. DATASETS AND ANALYSIS

In this section we introduce the datasets used in the study and perform an analysis of the distributions of temporal information in tweets which form the basis for SMERC.

A. Datasets

We collected tweets around the Australian Big Bash cricket league and Australian Rules Football League (AFL) based on predetermined hashtags for the events, using the Twitter streaming API. The hashtags used and number of tweets collected are outlined below. As with many sporting events, cricket matches and AFL games have specific hashtags, usually the names of the teams involved, that are publicized by the clubs and supporters prior to the event. While our method is for tweet clustering generally, such events are ideal for the present study as data collection can be planned in advance. Our collected data was stored in a MongoDB database and was later processed using custom Python 3.6 scripts.

To justify our temporal adjustment, we closely analyze the data from three Twitter datasets:

- Dataset D1: Tweets from the Australian Rules Football preliminary final match between the Adelaide Crows and the Geelong Cats on 22 September 2017.
Hashtag: #AFLCrowsCats
Number of tweets collected: 5018
- Dataset D2: Tweets from the 2017/18 Big Bash (the Australian men’s domestic Twenty20 cricket tournament) between Brisbane Heat and Melbourne Stars on 20 December 2017.
Hashtag: #BBL07
Number of tweets collected: 3153
- Dataset D3: Tweets from the 2017/18 Women’s Big Bash (the Australian women’s domestic Twenty20 cricket tournament) opening weekend on 9/10 December 2017.
Hashtag: #WBBL03
Number of tweets collected: 5393

We remark that similar results to the following were obtained across all sporting events studied. Our datasets are stored in a github repository¹.

B. Probability of tweets being related exhibits exponential decay with time difference

We manually labelled tweets from datasets D1 to D3 in order to examine how the probability of tweets being related varies depending on the time gap between them. As we will show, this decays exponentially, an observation which will form a critical part of the clustering algorithm SMERC we develop in Section III-A.

Our methodology to analyze the temporal relationship between related tweets is as follows:

- 1) Collect a series of tweets with a selected hashtag over a period of time.

¹<https://github.com/pete1729/phd-thesis>

- 2) Identify a set of on-field events (scoring events, penalties, etc), then manually cluster tweets that are in response to these events.
- 3) Record the time gaps between pairs of tweets where both tweets are in response to the same manually clustered event.
- 4) Record the time gaps between pairs of tweets, where only one of the tweets is in response to a manually clustered event, while the other is unrelated.
- 5) Bucketize the data and count the number of related pairs of tweets and unrelated pairs of tweets within each bucket.
- 6) Consequently measure the likelihood that a pair of tweets within a bucket is related, and examine how this varies with the time difference.

We first analyse dataset D1, the response to events in an AFL game. After manually classifying tweets, we have 18175 pairs of related tweets and 747779 pairs of unrelated tweets. We fit a Gaussian kernel to the associated time gaps in order to create a smooth density curve, implemented with the KernelDensity package from the sklearn.neighbors module in Python 3.6 [23]. The resultant density curves of time gaps between pairs of related and unrelated tweets are shown in Figure 1. As can be seen, the density of time gaps between unrelated tweets remains approximately constant with time difference, while the density of gaps between two related tweets about a given event decays over time. To determine the probability of tweets being related to the same event given the time separation, we bucketize the data and compute the proportion of related pairs of tweets to overall number of pairs within each bucket. We then plot this on a log-linear plot. As can be seen in Figure 2, we have roughly a straight line, corresponding to exponential decay. The curve has a slope of -0.01058 , giving that a suitable model for the probability that tweet pairs with gap time Δt are related is $Ce^{-0.01058\Delta t}$, where C is a constant.

In order to demonstrate that this property is not unique to AFL and instead holds across different sports, we repeat here the same analysis for the WBL dataset D3. For the WBL dataset, we have 7430 pairs of related tweets and 50431 pairs of unrelated tweets. As shown in Figure 3, we again have that the density of time gaps between unrelated tweets remains approximately constant with time difference, while the density of gaps between two related tweets about a given event decays over time. Bucketizing the data and plotting the probability curve in Figure 4 shows that a similar curve as fit to dataset D1 describes the data well. The fitted curve has slope -0.00692 , slightly less steep than for the AFL dataset D1. This indicates that for both datasets, the change in probability of tweets being related has an exponential relationship with the time gap between them. This is somewhat remarkable, due to the highly different nature of action for the two sports. Similar to baseball, cricket is a fundamentally ‘discrete’ game, whereas like soccer, AFL is fundamentally ‘continuous’. Note that a similar exponential

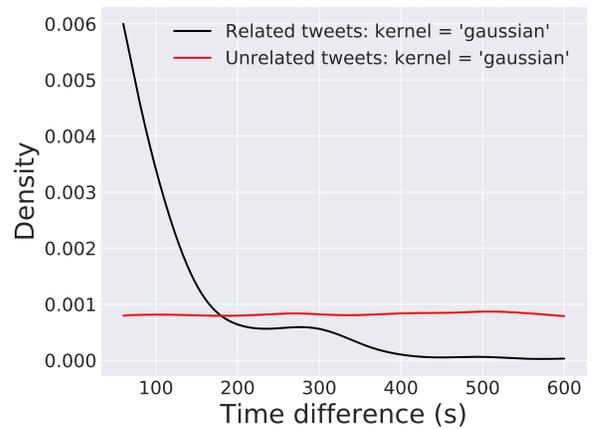


Fig. 1. AFL dataset D1: Density of time differences between pairs of related and unrelated tweets. The density of time differences between pairs of unrelated tweets stays roughly constant, while the density of time differences between related tweets decays. The slight bump at around 300 seconds is likely due to noise.

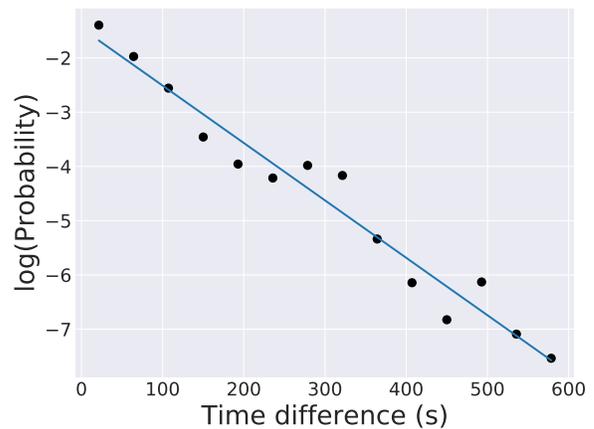


Fig. 2. AFL dataset D1: Log-linear plot of the probability of tweets being related, given a time separation. The straight line indicates exponential decay.

relationship was also found to hold in dataset D2, as well as for the other sporting events collected (not shown).

We have shown that the probability of tweets being related decays exponentially with the time gap between them. The exponential decay is likely explained by previous findings that human interest in topics decays exponentially over time [24]. This motivates a key new step in our clustering algorithm, outlined in Section III-A. To incorporate temporal information we multiply the textual affinity score between two tweets by a function exponentially decaying with the time gap.

III. CLUSTERING METHOD

A. SMERC algorithm

The clustering algorithm which we present here aims to create clusters of tweets in response to the same event. We do this in a several-stage process, where most of the steps

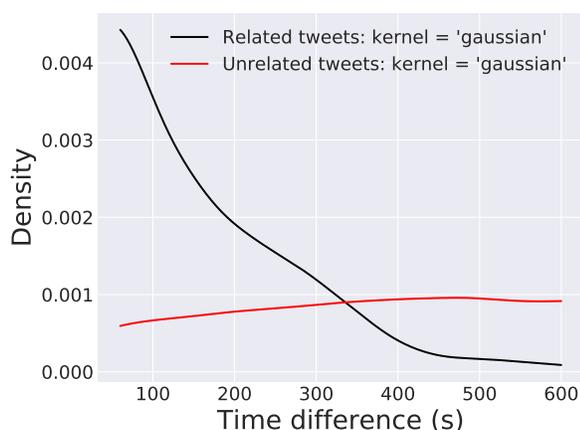


Fig. 3. WBBL dataset D3: Density of time differences between pairs of related and unrelated tweets. The density of time differences between pairs of unrelated tweets stays roughly constant with time difference, while the density of time differences between related tweets decays.

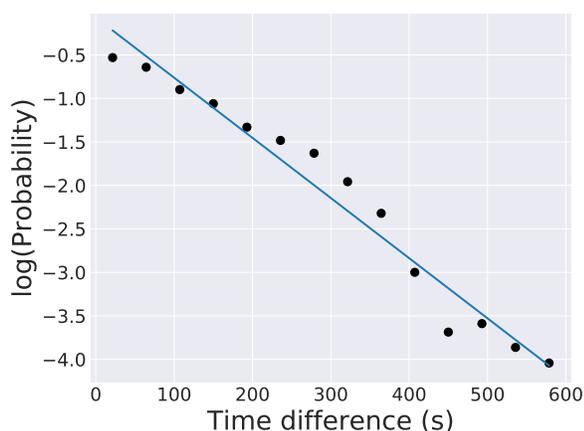


Fig. 4. WBBL dataset D3: Log-linear plot of the probability of tweets being related, given a time separation. The straight line indicates exponential decay.

are standard text processing and clustering. Our algorithm incorporates the novel use of temporal information in the form of an exponential decay function multiplied by the textual affinity between tweets. This prevents tweets that are a long time apart from being assigned to the same cluster. The clustering algorithm, including the text pre-processing steps, is summarized as follows:

- 1) Remove stop words from tweets.
- 2) Convert words to their associated stems.
- 3) Create a bag-of-words vector representing each tweet.
- 4) Use TF-IDF to more heavily weight words that occur less frequently.
- 5) Use cosine similarity to determine the textual similarity between TF-IDF vectors representing tweets.
- 6) Using the temporal information in each tweet, multiply by $e^{\Delta t_{ij}/T_p}$, an exponential decay function dependent on the time distance between tweets.
- 7) Use affinity propagation to determine clusters.

TABLE I
COMPARISON OF EVENT DETECTION PRECISION AND RECALL BETWEEN OUR METHOD AND OTHERS. NOTE THAT AS WE ARE TESTING ON DIFFERENT DATASETS, THESE RESULTS ARE INDICATIVE ONLY.

Method	Tweet topic	Precision	Recall	F1 score
SMERC	AFL and cricket	0.866	0.724	0.789
Twitinfo [26]	Soccer events	0.773	0.773	0.773
Twevent [16]	(Singaporean) news	0.762	0.619	0.683
MABED [27]	French politics	0.775	0.609	0.682

- 8) Filter clusters to ensure sufficiently high average affinity between elements.

Further detail on this algorithm is provided in the extended version of this paper [25].

Previous clustering methods are inferior at dealing with temporal information. Generally time is used as a linear variable which incorporates the loss of interest in topics over time less accurately. Another benefit of our method is that affinity propagation does not require selection of the number of clusters beforehand. It is entirely unsupervised and generates clusters of tweets related to events without requiring human assistance.

IV. EXPERIMENTS

We now test our proposed tweet clustering method and compare it to the following existing methods: Twevent [16], Twitinfo [26], and MABED [27]. We run our clustering algorithm SMERC on our three collected datasets. Through manually examining the cluster output, we determine whether each individual cluster is related to a real event. We then calculate the precision by dividing the number of clusters about real events, by the total number of clusters. An event is considered *missed* if our clustering algorithm conceivably could have detected the event, but failed to do so. We calculate the recall in the standard way, by dividing the number of true detected events by the total number of true events.

Due to the unavailability of the datasets used to test existing methods as well as usable code for these methods, we can neither test our method on previous authors' datasets, nor apply their methods to the datasets collected here. We will therefore reproduce the authors' precision and recall reported in their papers, and compare with the same values we calculate here. These values should therefore only be taken as indicative of the overall skill of each method. We will make the datasets used here, as well as code for SMERC, available upon acceptance of this paper.

A. Results

In Table I, we compare the precision and recall of our method to other event detection techniques. As standard test sets are not available, each method is evaluated on its own datasets. Despite this, and even though our method SMERC is tested on an arguably more challenging dataset where "events" are difficult to define, it performs remarkably well, showing the highest F1 score of all methods. SMERC also shows the

TABLE II
IMPROVEMENT TO CLUSTER QUALITY (THE PROPORTION OF TWEETS WITHIN A CLUSTER RELATED TO THE EVENT OF INTEREST) WITH TEMPORAL ADJUSTMENT. THIS SIGNIFICANT IMPROVEMENT IS THE KEY BENEFIT PROVIDED BY OUR METHOD.

Dataset	No temporal adjustment	With temporal adjustment
D1: AFL first prelim final	0.455	0.887
D2: Heat vs Stars BBL game	0.567	0.921
D3: WBBL first weekend	0.568	0.971

highest precision of all methods, and second-highest precision. We use the Twevent [16] approach for evaluating precision and recall, comparing the output clusters to real events. Of course, we note that as we are comparing on different categories of datasets, these results are indicative only. A better evaluation of our method is to investigate the quality of individual clusters, by looking at the types of tweets that are grouped together, particularly with the temporal adjustment.

However, our purpose in this paper is not merely a method that achieves high precision and recall statistics, but one which generates meaningful clusters about events. We therefore define the *quality* of a cluster as the proportion of tweets within a cluster which are related to the topic of the overall cluster. The improvement in the quality of the clustering from our temporal adjustment is determined by comparing average cluster quality both with and without the temporal adjustment.

Table II shows the increase in cluster quality after our temporal adjustment. This demonstrates the importance of the temporal adjustment as characterized by the exponentially-decaying model of user interest employed here. When applying this transformation to the affinity matrix A , the cluster quality increases substantially from around 50-60% to around 90-100% across all datasets. This is the primary benefit provided by our method: temporally distant tweets from unrelated events tend to not be mixed with clusters of tweets about a single event, and incorporating this exponential decay in interest improves the results dramatically. The extended version of this paper [25] explores the nature of the tweet clusters revealed by our method in greater detail.

As can be seen in Table III, the filtering of clusters to remove spam tweets and to ensure a sufficient average pairwise tweet affinity removes over half of the clusters. In particular, for the WBBL dataset which was collected over a weekend, time gaps between tweets are often higher. Consequently, a very high number of clusters are removed by the filtering.

The tradeoff between precision and recall can be controlled through the minimum average pairwise affinity between tweets. Setting a higher value for this parameter increases precision but reduces recall, and vice-versa.

V. CONCLUSIONS

In this paper we have developed a new method, *Social Media Event Response Clustering* (SMERC), for clustering tweets by using temporal information. This exponentially-decaying model was informed by a detailed analysis of a

number of Twitter datasets collected around sporting events, and is critical for improving the quality of clusters generated by the method.

In addition to sport, SMERC could be applied to other social science fields such as social unrest or natural disasters, where people respond to real world events. Also, in addition to Twitter data, this work could be applied to data from other social media platforms such as Facebook or Youtube. Future work will test SMERC to such data streams, as they become available.

At a computational level, there is also scope for future work in reducing the time complexity of the algorithm. SMERC could also be adapted into a system to detect events in real time, rather than using post processing. Furthermore, creating a dedicated dictionary for the topics that we are analysing would potentially allow the use of more modern tweet processing methods such as Tweet2Vec [28].

The difficulty of event detection varies depending on the topic of the event. Detecting rare events with known keywords such as earthquakes or goals in soccer, is much easier than detecting less well defined or frequently occurring events. After an earthquake, many people will tweet the word “earthquake”, a word which is rarely mentioned otherwise. Consequently, the choice of dataset will affect the measured performance of the algorithms.

Our clustering and time estimation approaches could potentially be improved by automated removal of noise on the input Twitter data [29], as done for other purposes [1]. Spam and advertisements tend to repeat identical or very similar tweets, which have high pairwise text similarity. Automated removal of these tweets before clustering will improve the amount of information content in the output clusters.

Unfortunately, obtaining code or datasets for tweet clustering and event detection methods is often prohibitive, which limits the ability to directly compare new methods to the existing state-of-the-art. This in turn makes determining what in fact constitutes “state-of-the-art” for a particular application effectively impossible. In the interests of reproducibility of work and having common tweet sets for testing, we make our datasets public and encourage other authors to do the same.

From a theoretical perspective, our work improves the understanding of the distribution of decay of interest reflected in online social media data streams around events. Practically, it provides an effective method to cluster tweets for the purpose of event detection. We evaluated this both quantitatively through the calculation of standard evaluation metrics such as precision and recall, but also qualitatively through inspection of the actual tweets clustered together by our method. We believe our method could be deployed by governments or other entities to conduct social sensing from microblogs.

REFERENCES

- [1] X. Liu, A. Nourbakhsh, Q. Li, S. Shah, R. Martin, and J. Duprey, “Reuters tracer: Toward automated news production using large scale social media data,” *CoRR*, vol. abs/1711.04068, 2017. [Online]. Available: <http://arxiv.org/abs/1711.04068>

TABLE III

DATASET CLUSTERING SUMMARY. FILTERING REFERS TO THE PROCESS OF REMOVING BOT TWEETS AND CLUSTERS WITH AN INSUFFICIENTLY HIGH AVERAGE PAIRWISE AFFINITY. THE WBBL DATASET HAS A HIGHER PERCENTAGE OF CLUSTERS REMOVED BY THIS MECHANISM AS IT WAS COLLECTED OVER A LONGER TIME PERIOD.

Dataset	# Collected tweets	# Clusters before filtering	# Clusters after filtering	# Clusters linked to specific events	# Missed events	Precision	Recall
D1: AFL first prelim final	5393	277	152	128	50	0.842	0.719
D2: Heat vs Stars BBL game	5018	430	160	144	45	0.900	0.762
D3: WBBL first weekend	3153	260	42	36	16	0.857	0.692

- [2] C. St Louis and G. Zorlu, "Can Twitter predict disease outbreaks?" *BMJ*, vol. 344, 2012. [Online]. Available: <http://www.bmj.com/content/344/bmj.e2353>
- [3] L. Mitchell and J. V. Ross, "A data-driven model for influenza transmission incorporating media effects," *Royal Society Open Science*, vol. 3, p. 160481, 2016.
- [4] R. van Noord, F. A. Kunneman, and A. van den Bosch, *Predicting Civil Unrest by Categorizing Dutch Twitter Events*. Cham: Springer International Publishing, 2017, pp. 3–16.
- [5] J. Tuke, A. Nguyen, M. Nasim, D. Mellor, A. Wickramasinghe, N. Bean, and L. Mitchell, "Pachinko Prediction: A Bayesian method for event prediction from social media data," *arXiv: 1809.08427*, 2018. [Online]. Available: <http://arxiv.org/abs/1809.08427>
- [6] T. Sakaki, M. Okazaki, and Y. Matsuo, "Earthquake shakes Twitter users: Real-time event detection by social sensors," in *In Proceedings of the Nineteenth International WWW Conference (WWW2010)*. ACM, 2010.
- [7] M. Hasan, M. A. Orgun, and R. Schwitter, "A survey on real-time event detection from the twitter data stream," *Journal of Information Science*, vol. 0, no. 0, p. 0165551517698564, 0. [Online]. Available: <https://doi.org/10.1177/0165551517698564>
- [8] J. Yang and J. Leskovec, "Patterns of temporal variation in online media," in *Proceedings of the Fourth ACM International Conference on Web Search and Data Mining*, ser. WSDM '11. New York, NY, USA: ACM, 2011, pp. 177–186. [Online]. Available: <http://doi.acm.org/10.1145/1935826.1935863>
- [9] M. R. Berthold and F. Höppner, "On clustering time series using euclidean distance and pearson correlation," *CoRR*, vol. abs/1601.02213, 2016. [Online]. Available: <http://arxiv.org/abs/1601.02213>
- [10] M. Gavrilo, D. Anguelov, P. Indyk, and R. Motwani, "Mining the stock market: Which measure is best," in *Proceedings of the sixth ACM SIGKDD international conference on knowledge discovery and data mining*, 2000.
- [11] D. J. Berndt and J. Clifford, "Using dynamic time warping to find patterns in time series," in *Proceedings of the 3rd International Conference on Knowledge Discovery and Data Mining*, ser. AAAIWS'94. AAAI Press, 1994, pp. 359–370. [Online]. Available: <http://dl.acm.org/citation.cfm?id=3000850.3000887>
- [12] E. Keogh and C. A. Ratanamahatana, "Exact indexing of dynamic time warping," *Knowledge and Information Systems*, vol. 7, no. 3, pp. 358–386, Mar 2005.
- [13] J. Macqueen, "Some methods for classification and analysis of multivariate observations," in *In 5-th Berkeley Symposium on Mathematical Statistics and Probability*, 1967, pp. 281–297.
- [14] J. Leskovec, L. Backstrom, and J. Kleinberg, "Meme-tracking and the Dynamics of the News Cycle," in *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM New York, NY, USA, 2009, pp. 497–506.
- [15] M. Gillani, M. U. Ilyas, S. Saleh, J. S. Alowibdi, N. Aljohani, and F. S. Alotaibi, "Post summarization of microblogs of sporting events," in *Proceedings of the 26th International Conference on World Wide Web Companion*, ser. WWW '17 Companion. Republic and Canton of Geneva, Switzerland: International World Wide Web Conferences Steering Committee, 2017, pp. 59–68. [Online]. Available: <https://doi.org/10.1145/3041021.3054146>
- [16] C. Li, A. Sun, and A. Datta, "Twevent: Segment-based event detection from tweets," in *Proceedings of the 21st ACM International Conference on Information and Knowledge Management*, ser. CIKM '12. New York, NY, USA: ACM, 2012, pp. 155–164. [Online]. Available: <http://doi.acm.org/10.1145/2396761.2396785>
- [17] S. Zhao, L. Zhong, J. Wickramasuriya, and V. Vasudevan, "Human as real-time sensors of social and physical events: A case study of Twitter and sports games," *CoRR*, vol. abs/1106.4300, 2011. [Online]. Available: <http://arxiv.org/abs/1106.4300>
- [18] Y. Takeichi, K. Sasahara, R. Suzuki, and T. Arita, "Concurrent bursty behavior of social sensors in sporting events," *PLOS ONE*, vol. 10, no. 12, pp. 1–13, 12 2015.
- [19] Q. Zhao, M. A. Erdogdu, H. Y. He, A. Rajaraman, and J. Leskovec, "Seismic: A self-exciting point process model for predicting tweet popularity," in *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ser. KDD '15. New York, NY, USA: ACM, 2015, pp. 1513–1522. [Online]. Available: <http://doi.acm.org/10.1145/2783258.2783401>
- [20] Y. Zhu, E. Zhong, S. J. Pan, X. Wang, M. Zhou, and Q. Yang, "Predicting user activity level in social networks," in *Proceedings of the 22nd ACM International Conference on Information & Knowledge Management*, ser. CIKM '13. New York, NY, USA: ACM, 2013, pp. 159–168. [Online]. Available: <http://doi.acm.org/10.1145/2505515.2505518>
- [21] Y. Gandica, J. Carvalho, F. Sampaio dos Aidos, R. Lambiotte, and T. Carletti, "Stationarity of the inter-event power-law distributions," *PLoS ONE*, vol. 12, no. 3, pp. 1 – 10, 2017.
- [22] P. Mathews, L. Mitchell, G. Nguyen, and N. Bean, "The Nature and Origin of Heavy Tails in Retweet Activity," in *The 26th International Conference on World Wide Web Companion*, 2017, pp. 1493–1498.
- [23] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [24] L. Li, L. Zheng, F. Yang, and T. Li, "Modeling and broadening temporal user interest in personalized news recommendation," *Expert Syst. Appl.*, vol. 41, no. 7, pp. 3168–3177, Jun. 2014. [Online]. Available: <http://dx.doi.org/10.1016/j.eswa.2013.11.020>
- [25] P. Mathews, C. Gray, L. Mitchell, G. T. Nguyen, and N. G. Bean, "SMERC: Social media event response clustering using textual and temporal information," *arXiv: 1811.05063*, 2018. [Online]. Available: <https://arxiv.org/abs/1811.05063>
- [26] A. Marcus, M. Bernstein, O. Badar, D. Karger, S. Madden, and R. Miller, "Twitinfo: Aggregating and visualizing microblogs for event exploration," in *Proceedings of the 2011 annual conference on Human factors in computing systems*. ACM, 2011, pp. 227–236.
- [27] A. Guille and C. Favre, "Event detection, tracking, and visualization in twitter: a mention-anomaly-based approach," *CoRR*, vol. abs/1505.05657, 2015. [Online]. Available: <http://arxiv.org/abs/1505.05657>
- [28] B. Dhingra, Z. Zhou, D. Fitzpatrick, M. Muehl, and W. Cohen, "Tweet2vec: Character-based distributed representations for social media," in *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Association for Computational Linguistics, 2016, pp. 269–274. [Online]. Available: <http://www.aclweb.org/anthology/P16-2044>
- [29] M. Nasim, A. Nguyen, N. Lothian, R. Cope, and L. Mitchell, "Real-time detection of content polluters in partially observable Twitter networks," in *Proceedings of the 26th International Conference on the World Wide Web (WWW '18) Companion*, 2018, pp. 1331–1339.