

Personal Life Event Detection from Social Media

Smitashree Choudhury
Knowledge Media Institute
The Open University
United Kingdom
smitashree.choudhury@open.ac.uk

Harith Alani
Knowledge Media Institute
The Open University
United Kingdom
h.alani@open.ac.uk

ABSTRACT

Creating video clips out of personal content from social media is on the rise. MuseumOfMe, Facebook Lookback, and Google Awesome are some popular examples. One core challenge to the creation of such life summaries is the identification of personal events, and their time frame. Such videos can greatly benefit from automatically distinguishing between social media content that is about someone's own wedding from that week, to an old wedding, or to that of a friend. In this paper, we describe our approach for identifying a number of common personal life events from social media content (in this paper we have used Twitter for our test), using multiple feature-based classifiers. Results show that combination of linguistic and social interaction features increases overall classification accuracy of most of the events while some events are relatively more difficult than others (e.g. new born with mean precision of .6 from all three models).

Keywords

Social Web, social media, event detection, personal life events

1. INTRODUCTION

With the wide spread of social media sites (e.g. Twitter, Facebook, YouTube), millions of people use them on daily basis to communicate and share information on a wide variety of events, ranging from world events (e.g. World Cup), to personal events (e.g., Wedding, Graduation). Use of these systems serves the multitude of purposes of knowledge sharing, information communication, event organisation, professional collaboration, political expression, as well as socialisation. To put in perspective, more than 500 million of tweets generated in a day¹, millions of photos are uploaded to Facebook every day. There may be differences in terms of content volume created on different platforms depending on the personal preferences and the perceived purpose of the

¹<https://blog.twitter.com/2013/new-tweets-per-second-record-and-how>

tool, nonetheless most popular online systems are carrying huge amount of data created by individual users in the form of texts, videos, and photos. While technology for data creation and storage has significantly matured and efficiently managed, accessing, managing and processing of such data is still a challenge and can be done by few experts. Due to the lack of efficient data access mechanism available to normal users, most of the historical data tend to be forgotten or will remain unused.

Access and reuse of such information trove will provide greater insight about the individual user, their preferences, and situational dynamics and result in many useful applications e.g. personalised healthcare, customised training and education, social and community engagement application and life stories. To this end, mining and analysing such content could help identifying one's life milestones and salient events. Identifying interesting and important moments in one's timeline on social media is valuable to services such as Facebook Lookback and Google Awesome, which generates short video clips for users to summarise and visualise their timelines.

In realisation of the importance of events on social media, Facebook² has recently generated millions of 1 minute look-back videos of content from users' timelines. Over 270 million video rendered and over 200 million users watched their look back movie in the first two days and more than 50% shared their movie. A project like Intel's Museum of Me³ follows a similar line to collect data from user's Facebook profile and generate a short video. Purpose of our work (personal life event detection) is a sub-objective of the broader research objective in similar direction i.e, automatic creation of digital documentaries from social media content including interesting and relevant life moments and events.

Event detection from social media content has so far been focused on detecting world events such as earthquakes [Chile, japan], political protests, elections (US, Germany, UK) and planned public events such as entertainment award functions (Oscar, Golden Globe), academic events (conferences), sports event (Olympic). However, detection of personal life events have been mostly overlooked, and only mildly investigated for content recommendation [cite]. Objective of this piece of is to automatically identify interesting and impor-

²<https://code.facebook.com/posts/236248456565933/looking-back-on-look-back-videos>

³<http://www.intel.com/museumofme/r/index.htm>

tant life events of individual users from their social media content, which can be part of their personal digital storybook or memory archive. In this work, we have taken Twitter as the test platform and will extend our research to other systems such as Facebook, Instagram, Pinterest in our future work.

Detecting personal events is non-trivial and may require a combination of multiple approaches for a robust detection result. Unlike public events or events concerning celebrities and well-known personalities, personal events may not be characterised by high activity volume and additional sources of information e.g. blogs or Wikipedia. These events are limited to the concerned person and to her immediate social network (friends and family). In addition to the above problems, microblog sites like Twitter bring its own complexities with short, informal and noisy content. Any meaning-making task on these content has to deal with these idiosyncrasies. Next, we will briefly delve into the concept of a personal event before going into the details of the experimental work.

1.1 Personal Life Events

Personal life events range from recurring events such as birthdays and anniversaries, to very occasional and uncommon events, such as work promotions, and relocation. Events can also be further categorised on an affective scale, from highly positive and pleasant events to unpleasant events, such as illnesses or accidents and deaths of loved ones. In this paper, we focus on 5 life events (4 positive and 1 negative) i.e. graduation, marriage/engagement, new job, birth of child, and surgery. Our motivation to start with these events inspired by a study [?] which lists 6 important memorable life events are "Beginning school", "first full time job", "Falling in love", "Marriage", "Having children;", "Parent's death".

The main contributions of this paper are not on algorithm and its efficiency, but rather on presenting evidence that with effective combination of existing methods and social media data, we can analyse and detects important and critical moments of individuals life., hence the contributions are:

- a thorough study of five personal life events and their idiosyncrasies as reported in social media especially in Twitter .
- detection of life events using both content and interaction features.

This paper is organised as follows: In section 2 we review related work in the field of event detection in social media and in section three, we briefly describe how personal life events are reported on twitter and their characterisation. Section 4 describes our approach which includes feature selection and model construction followed by discussion and conclusion in section 5.

2. RELATED WORK

Event detection is now a new research subject, and has been part of studies on topic detection in news stories and other text documents [?]. Social media brought multi modal content created by both professional and amateurs leading to a

resurgence of interest in detecting social topics and events in this new domain[?]. We have been motivated by the need to identify personal life events, which have a great personal value when aggregated over time and location. One of the prerequisites of such a system is the identification of content reporting a real event. Events can be planned events such as cultural events, tech conferences, music award functions, elections or sports event or unplanned events for example, natural disasters, earthquake [?] and even generic events such as breaking news events are subject of few studies [?][?]. Existing studies cover both planned and unplanned events with varying degrees using both machine learning and text analysis techniques. Benson et.al.[?] reported detecting concert events from social media stream using city calendar as a target list. Agarwal et. al.[?] detected events such as factory fire, labor strike from Twitter stream using a combination of local sensitive hashing and location dictionary. Weng and Lee[?] proposed event detection with clustering of word bursts from tweets. Authors in [?] proposed a natural disaster alert system using Twitter users as virtual sensors. In their work, they were able to calculate the epicentre of an earthquake by analyzing the delays of the first messages reporting the shock. Social media centric event detection also covers non textual data such as photos and videos, Chen et al.[?] discovered social event from Flickr photos by using both user tags and other metadata including time and location (latitude and longitude). Firan et.al[?] explored tags, title and description to classify pictures into event categories. Some of the popular approaches used for event detection are spatio-temporal segmentation[?], burst analysis in word signals, clustering as well as topic detection techniques.

To the best of our knowledge, we found no prior studies on personal life event detection from social media except one reported in [?] where authors tried to detect two life events "marriage" and "employment" and bears some similarity to our work. Our focus is on user level event detection that can be used to build individual digital storyboards from historical data.

3. PERSONAL EVENTS ON TWITTER

We now define the concept of personal life event in the context of Twitter message stream and provide a definition of the problem that we address in this work.

Definition of term "event" differs from domain to domain ranging from Philosophy to cognitive psychology to computing. Despite a lack of uniform definition of the term it embeds a few generic characteristics such as time, participating objects and a location. In this context, we define an event as a real world occurrence with an associated time period and one or more participating objects/agents at a certain location which may or may not be explicitly apparent in tweet messages. According to this definition a tweet needs to reflect a time interval when the event has occurred involving either the user or someone connecting to the user as the participating agent. Based on this abstract notion, we looked into the real data to confirm or re-arrange the definition and devise a strategy for detecting personal events.

3.1 Dataset

As a first step, we collected tweets using Twitter streaming API⁴ which allows to crawl some portion of public tweets as and when it comes. We restricted tweets to English language only and crawled for 3-4 hours per day for three weeks. The entire dataset contained around 4 million tweets. Ratio of event tweets to non-event tweets is expected to be extremely skewed as the targeted events are very specific and user centric. So the next logical step is to use a filter mechanism to segregate the event related tweets from the rest and process further. For this initial segregation, we extended the event query with synonyms and related terms and phrases (shown in Table 1). These related terms are mainly synonyms and terms commonly known and used to describe the event of interest. Use of related terms with the main event terms were intended to widen the coverage where users might not be using the exact terms to describe the main events. After filtering we got 9168 tweets for marriage event, 2570 tweets for graduation, 3192 tweets for surgery, 3661 for new job and 2954 tweets for new born. A question may arise about those tweets where the event term may be absent yet the implicit semantics reflects a real event for example. "Welcome to the new member of our family". However, we agree such kind of possible omissions with the present approach and intend to capture them with contextual and historical information as part of our future work. The resulting filtered datasets still contain many irrelevant tweets. For example, "family have brought a 2nd lawsuit against her, this time to try to annul her marriage" is not about a marriage event though it contains the keyword. Our task is identify such tweets from genuine event tweets by means of binary classification.

Table 1: Events are their related words.

Event terms	Related Terms
Marriage	"Wedding", "Tied the knot", "married"
Graduation	"Convocation", "commencement "
New Job	" new position", "first day at work", "job offer"
New Born	"Baby boy", "baby girl", "new born"
Surgery	"Operation"

Manual inspection of these tweets revealed that event reporting tends to happen at three time spans; part, present, and future. We also noticed three categories of participating agents (self, others individual and general public). Examples of such diversities are shown in table 2.

In light of these findings, defining a personal event seems to be more tricky and imprecise. Two pertinent questions here are how to resolve the time reference associated with the event and how to associate the right subject (participating agent) with the event. In this study we are only focusing on the events where the time reference can be resolved to a specific time point within a month time interval by automatic means. One such example is "I graduated yesterday", " 26 days to graduation". In both cases, the time of the event can be resolved with help from the timestamp attached to the message. However, ambiguous time references such as "graduation is so close yet so far", "marriage in few weeks time" are ignored.

⁴<https://dev.twitter.com/docs/api/streaming>

The second dimension where the event reporting differs is on participating agent or affected subject. Event tweets are either about the user who created the tweet or about someone else known to the user and in some cases, about an undefined group of people e.g. group of students. Since our focus is on personal events, ideally we should target self-reported tweets and ignore the rest. But resolving an event to a participating agent needs advanced semantic role labelling which will be our next step of this ongoing work. For this paper, we restricted our attention to generic event detection, hence included all the tweets irrespective of who the affected subject is.

Based on this generic definition, we proceed with our actual experiment task that starts with feature extraction.

4. FEATURE EXTRACTION

After filtering event related tweets from the non-event tweets, we extracted different types of features [?] to be used for building event classifiers. We examined several feature categories describing different aspects of tweets and users. Specifically we considered lexical, sentimental and social interaction features.

4.1 Textual Features

Event term: The basic lexical feature of an event is the event term itself and most closely related terms or its synonym "#graduation, convocation" for the event graduation. The synonyms are extracted from Wordnet⁵

Co-occurring textual Features are the features of a term that co-occur significantly along with the event term for example, "cap", "dress", "present", "prom", "party" are some of the frequently occurred terms for graduation, while "prayer", "hospital" for surgery. Presence of these terms along with the main event term is expected to boost the detection process. Co-occurring terms were extracted from various tag based social media sites such as Flickr, instagram where terms are described with highly related terms. These features are event specific and treated as binary values i.e. 1 for presence otherwise 0.

Temporal terms: This feature reflects the presence of time terms in a tweet. Since the content are about an event, it is intuitive to assume that some reference to time is natural and required by definition. For this feature, we used LIWC's time category which includes 68 time terms.

Person reference terms: Since these events are about personal life event one or more reference terms reflecting social relation is expected when the event is about somebody other than the poster, or self reference if the event is about the user.

Sentiment: personal events are expressed with rich emotions both for pleasant or unpleasant events. Sentiments are detected by Sentistrength [?] library and proved to be good for social media sentiment detection. Value of this feature ranges from -5(negative) to +5(positive) while +1 to -1 considered as neutral.

⁵<http://wordnet.princeton.edu/>

Table 2: Events and their examples from Twitter.

Event	Examples
Marriage	Kansas City here we come! It’s happening! My sister’s marriage this weekend!! :) 8 years ago this day , married to the most loving man on this earth. Congratulations to my beautiful friend, @SheridanMills, who tied the knot today! ???
Graduation	Happy graduation day, bebe! Congrats cutie pie! http://t.co/YqgNgK9WMw Graduation is just around the corner. Time to start planning programs and certificates. Talk to our print consultants today! 3 sets of graduation picture next week! Hahaha. At last! :)
New Job	First day of a new job.... Kind of dreading it. #officeassistant Starting my new position today. Ayy lmao. Shout out to my cuz Quincy Johnson aka Q. On his new Executive Chef position! ???
New Born	My baby girl is here! Introducing: Halen born naturally May 3rd @ 4:43 pm. Exactly 3 weeks till my babyshower & almost 7 weeks till my baby boy Is born ?
Surgery	Good luck on your surgery today @chloebieber ear surgery ??it went well Everyone please continue to pray for Karlie these next 5 hours. She just went back for her brain surgery. #PrayersForKarlie

Non-Textual and punctuation Features relating to punctuation and emoticons such as presence of "!"/?" are expected to add the discriminating qualities of a learning model.

4.2 Interaction and Social Feature

Unigram is a basic model for classification and the result shows a reasonable accuracy including a poor performance for the *new born* event. This motivated us to further explore the feature space and extract more defining attributes of an event in terms of activity and interactions based on the simple logic that important events are bound to generate more attention and activity within the immediate personal network of an individual. Accordingly, we computed the following Twitter specific features concerning to a tweet and the user. These features can be broadly classified into two categories: **1) Activity and 2) Attention**. Activity features (first four in the list below) are based on user’s activity (tweets, re-tweet and replies) while attention features are the measures of engagement between the user and his/her network (last four features in the list below)

1. Tweets per day: Number of tweets per day a user posts
2. Re-tweets per day: Number of tweets per day a user posts.
3. Replies per day: Number of replies given by the user to other users.
4. Unique mentions per day: Number of unique mention (users addressed) in a day by the user.
5. Number of times the user is mentioned in a day
6. Number of times a user is replied to, by other users
7. Number of times a tweet is re-tweeted by other users
**
8. Number of times a tweet is marked as "favourite" by other users.**

In this work, we have used the last two interaction features only for comparison study, while other features are part of an extension work primarily focusing on iteration specific models in identifying life events.

5. EXPERIMENTAL RESULT

In this step, we analyse the experimental steps and present the results of classifications. We started with the ground-truth annotation process followed by classification steps and their results.

5.1 Ground Truth Annotation

In the absence of any benchmark data for personal event detection prepared a gold standard dataset with manual annotation of 2 users with computing background . Annotators were given 1000 tweets per event for annotation. These 1000 tweets are randomly selected from the filtered dataset. Instruction for annotation was to annotate a tweet as event positive (presence of event) if they consider the tweet describes an event happening (present e.g. today) or about to happen with certainty (e.g. 4 days to graduation) within a month’s time window. It is difficult to precisely define an event as most of the tweets are not reported exactly during the event but pre and post event. Since our objective is to identify the event from user’s timeline with definitive time stamp attached to the event, we opted for a 1 month time interval. We retained those tweets (304) as event positive tweets whenever both the annotators agreed on the label. It is imperative to mention that event negative tweets are simply those where annotators felt that a particular event is not occurring despite the presence of event related keyword.

5.2 Event Detection: Unigram Model(UNI)

Our first model is the simplest bag-of-word model where word frequencies are used as features for document classification. In our case, each tweet is considered 1 document. We first applied a String to word vector filter that converts the strings into numerical features. Then we trained our model with 10-fold cross validation using four different types of classifiers: Naive Bayes (NB), Multinomial Naive Bayes (MNB), Support Vector Machine (SVM) and Decision Tree

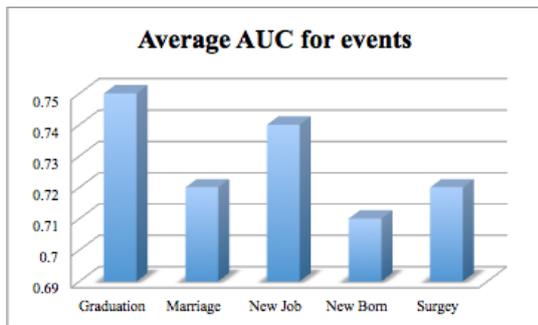


Figure 1: AUC curve for different events.

(J48) implemented in machine learning library Weka [?]. We evaluated our model on the test set (100 from each event) and performance of these classifiers reported in terms of Recall (is the number of correct results divided by the number of results that should have been returned) Precision (is the number of correct results divided by the number of all returned results) and F-score (harmonic mean). Table 3 (fig. 2) shows the average precision, recall and F score for all the events. However SVM performed best in 4 out of 5 followed by Naive Bayes. Graduation (.8) has highest precision score whereas "New job" has the highest recall (.95) score. The most difficult event is the "New born" across all the classifiers with lowest precision score (.55).

Examining the ROC curves which plots the true positives (TP) vs false positives (FP) and indicates the area under curve (figure 1) (AUC: probability that a classifier will rank a randomly chosen positive instance higher than a randomly chosen negative example) ranges from .71 to .75 giving a reasonable quality of the learners. NB performs better than SVM with an average of .77 against .72 across all events.

Table 3: Average precision, recall and f-Measure from all classifiers based on unigram model.

Event	Precision	Recall	F-Measure
Graduation	0.80	0.80	0.73
Marriage	0.75	0.87	0.79
New Job	0.78	0.95	0.80
New Born	0.55	0.92	0.68
Surgery	0.72	0.87	0.76

Analysis of error classification mainly showed the diversity of language constructs among the misclassified tweets. Since the model is purely content based, any variation not captured by the model are missed from the result.

5.3 Event Detection: Model with Contextual Lexical Patterns (UNI+META)

Bag-of-words or unigram model is the basic approach yet proved to have reasonable accuracy though with lots of false positives. This led us to refine the model with more lexical features and features such as sentiment. We considered features (described in sec. 4) such as co-occurring terms (e.g. prayers, hospital for surgery), POS tagging, presence of so-

cial relation terms (my friend, sister etc.), temporal terms (today, week, morning etc.), sentiment strength of a tweet. POS tagging was done using Stanford tagger⁶ and sentiment was derived using the Sentistrength java library[?].

Recognizing Temporal Expression: Temporal features tend to be implicit, diverse, and informal (e.g. last week, hourly, around the corner). Identifying these references within the vicinity of an event term occurrence increases the likelihood of accurate detection. Moreover, we need to resolve the tense of the verb as well to know whether the tweet is about some future event, or past. In this paper, we are using the time terms of LIWC dictionary which has 68 time inducing terms (e.g. forever, week, until etc.). This feature also used as a binary feature in the second classification model.

Average accuracy of the second model showed an average improvement of 4-5 % in precision score over the initial model for all the events, showing that simple lexical features are able to capture some of the diversity. For brevity purpose we are only showing the results of the top classifier (SVM).

Table 4: Precision, Recall and F-measure for (UNI+META) Model (SVM).

Event	Precision	Recall	F-Measure
Graduation	0.83	0.81	0.819
Marriage	0.77	0.83	0.798
New Job	0.818	0.93	0.865
New Born	0.61	0.92	0.733
Surgery	0.77	0.87	0.816

5.4 Event Detection: Model with Interaction Features (UNI+META+INT)

Inherent in social media and social networks, it is intuitive to hypothesize that interesting events will stimulate interesting and increased interaction among the friend circle of the user in the form of replies and sharing. The third and the final model takes advantage of these interaction features embedded in microblogging sites through mechanisms like retweet and favourites. Each tweet is now represented with two more features besides the above lexical features for classification. We used only SVM as the classifier because of its superior performance in previous two occasions. Results of the final model (table 5) are reported by means of precision score per event. A final comparison of four models (UNI, UNI+META, UNI+META+INT and INT) is shown in figure 3. The result shows that, although the hybrid model performed better than the unigram-based one (UNI), the improvement was marginal. On the other hand, the model based only on interaction features (INT) performed worst, where accuracy dropped to 53-61%.

6. CONCLUSION

This paper describes event detection from personal timeline of a user in Twitter. Existing detection tasks predominantly focused on public events and events concerning celebrities both from news articles and social media whereas personal life events are mostly overlooked. We started with 5 life

⁶<http://nlp.stanford.edu/software/tagger.shtml>

Table 5: Precision, Recall and F-measure for (UNI+META+INT) Model (SVM).

Event	Precision	Recall	F-Measure
Graduation	0.85	0.83	0.839
Marriage	0.79	0.83	0.809
New Job	0.82	0.91	0.862
New Born	0.64	0.92	0.754
Surgery	0.78	0.87	0.822

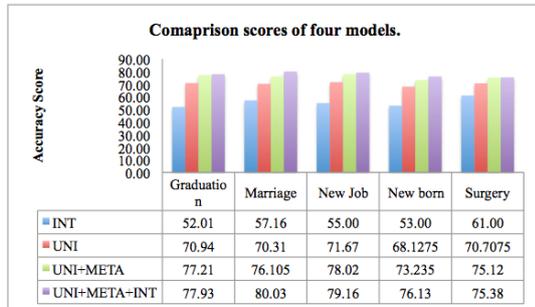


Figure 2: A comparative performance of four different models.

events and trained 5 different binary classifiers based on bag-of-word features which gave 55 to 80% precision on a test dataset with an average AUC of 77%. The learning models were further streamlined with meta features such as sentiment, temporal, social relation terms, emoticons and punctuations features, which improved the classification performance by 4-5%, however addition of interaction feature in the third classifier did not yield substantial improvement contrary to the expectation. This final result is a stronger motivation for an in-depth analysis of these features in our future work. We also aimed to adopt an unsupervised approach to detect life events as there may be many more unexpected events happening in one's life bearing substantial influence in life and eligible to be included .

7. ACKNOWLEDGMENT

This work was supported by EPSRC project ReelLives (EP/L004062/1).

8. REFERENCES

- [1] P. Agarwal, R. Vaithyanathan, S. Sharma, and G. Shroff. Catching the Long-Tail : Extracting Local News Events from Twitter. In *book1*, pages 379–382, 2012.
- [2] E. Benson, A. Haghighi, and R. Barzilay. Event Discovery in Social Media Feeds. In *book1*, volume 3, pages 389–398. Association for Computational Linguistics, 2011.
- [3] L. Chen and A. Roy. Event detection from flickr data through wavelet-based spatial analysis. In *Proceedings of the 18th ACM Conference on Information and Knowledge Management*, CIKM '09, pages 523–532, New York, NY, USA, 2009. ACM.
- [4] B. D. Eugenio, N. Green, and R. Subba. Detecting Life Events in Feeds from Twitter. pages 274–277. Ieee, 2013.
- [5] C. S. Firan, M. Georgescu, W. Nejdl, and R. Paiu. Bringing order to your photos: Event-driven classification of flickr images based on social knowledge. In *Proceedings of the 19th ACM International Conference on Information and Knowledge Management*, CIKM '10, pages 189–198, New York, NY, USA, 2010. ACM.
- [6] J. Glšck and S. Bluck. *Looking back across the life span: A life story account of the reminiscence bump*. Springer, 2007.
- [7] Q. He, K. Chang, and E.-P. Lim. Analyzing feature trajectories for event detection. In *Proceedings of the 30th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '07, pages 207–214, New York, NY, USA, 2007. ACM.
- [8] A. Jackoway, H. Samet, and J. Sankaranarayanan. Identification of live news events using Twitter. In *book1*, page 1, New York, New York, USA, 2011. ACM Press.
- [9] A. Java, X. Song, T. Finin, and B. Tseng. Why we twitter: Understanding microblogging usage and communities. In *Proceedings of the 9th WebKDD and 1st SNA-KDD 2007 Workshop on Web Mining and Social Network Analysis*, WebKDD/SNA-KDD '07, pages 56–65, New York, NY, USA, 2007. ACM.
- [10] S. Papadopoulos, C. Zigkolis, Y. Kompatsiaris, and A. Vakali. Cluster-based landmark and event detection for tagged photo collections. In *book1*, volume 18, pages 52–63, Los Alamitos, CA, USA, Jan. 2011. IEEE Computer Society Press.
- [11] S. Phuvipadawat and T. Murata. Breaking news detection and tracking in twitter. In *Proceedings of the 2010 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology - Volume 03*, WI-IAT '10, pages 120–123, Washington, DC, USA, 2010. IEEE Computer Society.
- [12] T. Sakaki. Earthquake shakes twitter users : Real-time event detection by social sensors. In *Proceedings of the 19th International Conference on World Wide Web*, 2009.
- [13] M. Thelwall, K. Buckley, G. Paltoglou, and D. Cai. Sentiment strength detection in short informal text, 2010.
- [14] C. L. Wayne. Topic detection tracking (tdt). In *In Proceedings DARPA Broadcast News Transcription and Understanding Workshop*, page 98, 1998.
- [15] J. Weng, Y. Yao, E. Leonardi, F. Lee, and B.-s. Lee. Event detection in twitter. In *book1*, pages 401–408. Ieee, 2011.