2017 IEEE 11th International Conference on Semantic Computing

Using Social Sensing to Discover Trends in Public Emotion

Maryam Hasan, Elke Rundensteiner, Xiangnan Kong, and Emmanuel Agu
Computer Science Department
Worcester Polytechnic Institute
100 Institute Rd, Worcester, MA, USA
Email: mhasan@wpi.edu, rundenst@cs.wpi.edu, xkong@cs.wpi.edu, emmanuel@cs.wpi.edu

***Abstract*—The rapid growth of social media, such as twitter, provides a great opportunity for identifying and analyzing people's emotions in response to various public events, such as epidemics, terrorist attacks and political elections. Detecting the emotions of people on different events are crucial in many applications. However, the high volume and fast pace of social media make it challenging to analyze public emotions from social media data in real-time. In this paper we propose a method to measure public emotion and predict important moments during particular public events. Given a stream of tweets, we analyze the impact of major public events, both tragic and enthusiastic ones, on public emotion. We develop a full-stack architecture that performs real-time emotion analysis on Twitter streams. We design a supervised learning approach for classifying tweets based on the type of the emotion they elicit. Then we aggregate each emotion class to discover emotion-evolving patterns over time. We also propose an online approach to predict emotion-intensive moments during real-life events. Our emotion analysis methodology is shown to present a fast and robust way of analyzing online stream of tweets.**

I. Introduction

Social networks and microblogging tools such as Twitter are increasingly used by individuals to express their personal feelings and opinions in the form of short text messages [1]. These text messages may contain emotional indicators of individuals such as happiness, anxiety, and depression. In fact, social networks such as Twitter often contain a rich combination of emotions [1], which makes them appropriate data sources for behavioral studies, especially for studying emotions of individuals, as well as the general public.

Moreover, social networks have been used as one of the prevalent communication channels for spreading news [2]. They allow users to share their opinions and feelings about real-time events as they occur. People may write posts about local events such as sport games, political elections, or social issues. For example, the disappearance of the airplane from Malaysia Airlines in 2014 and the Ebola virus epidemic burst in West Africa in late 2013 were extensively reported by Twitter users. Analyzing these events can provide valuable information about reactions and emotions of people regarding the events which can not be achieved using traditional media. The growth of social networks such as Twitter now empowers us to identify the influence of a social event on a large group of people in near real-time.

Population level studies of emotions could be beneficial in a variety of fields including social science, political science, public health research, and market research which are interested in aggregate emotion, instead of individual cases. It could assist government agencies in recognizing growing public fear or anger associated with a particular decision or event, or in helping them to understand the public's emotional response toward controversial issues or international affairs. In some cases rapidly gaining such insights as well as getting a deeper understanding on trends associated with positive versus negative emotion propagation across a population can be critical. The public emotion analysis can aid public health researchers by providing (1) a low-cost method to measure potential risk across different sub populations; (2) useful knowledge for identifying at-risk populations; and (3) a method to help formulate new hypotheses about the impact of real-time events on populations.

In this paper, we focus on studying public emotion through analyzing emotion trends driven by external public events. We select two different classes of events, negative tragic events as well as positive events. We analyze the impact that major real-life events tend to have on the public emotion.

To detect the emotion expressed in text messages, we have developed a supervised learning system called Emotex to automatically classify the text messages of users into their emotional states [3]. Using Emotex in a controlled environment on curated data, we were able to train supervised classifiers with up to 90% prediction accuracy.

In this work we now aim to deploy the trained model in the wild to analyze real-life events. For this purpose, We first apply our Emotex system [3], [4] to automatically detect the emotion of people from their text messages. It first learns an emotion classification model from a large dataset of emotion-labeled messages. The emotion classification model is created using curated data in a controlled environment, where data is filtered and preprocessed. Then, the model is deployed in a two-stage framework to classify the raw streams of tweets posted about an event. A binary classifier is created in the first stage to separate tweets with explicit emotion from tweets without emotion. The second stage utilizes our emotion classification model for a fine-grained emotion classification of tweets with explicit emotion. Then we aggregate text messages at the emotion class level to analyze public emotion trends driven

978-1-5090-4284-5/17 $31.00 © 2017 IEEE
DOI 10.1109/ICSC.2017.76

by social events and discover emotion-evolving patterns over time. We also propose an online approach to measure public emotion and predict emotion-intensive moments during social events. Our approach is able to analyze the real-life events, in spite of the noise, linguistic diversity, and fast-evolving nature of tweets in the wild. We deploy our approach in the wild to analyze public emotion trends during different types of real-life events.

Previous work on detecting important moments or sub-events during an event rely on the frequency of similar tweets [5]–[10]. However we are looking for the emotion-intensive moments with high impact on public emotion. Thus, we predict important moments based on the number of tweets within a specific emotion class, instead of the frequency of similar messages. More precisely we aim to find temporal bursts of public emotion during real-life events. Such temporal bursts of the emotional content of tweets can point towards important moments. Detecting important moments of an event is critical as they can provide a summary about reactions and emotions of people regarding the event and give an overall insight about the event. For example, in a soccer game users are likely to reflect their emotion during important moments such as goals and red cards.

We examine our emotion analysis approach using a large crawl of Twitter stream data about real-life events. Our emotion analysis system is shown to present a fast and robust way of analyzing tweet streams in the wild. It provides insights about how people feel about important events on social networks such as Twitter. Our experiments investigate the research questions highlighted below:

- Measure the impact of real-life events on people: how the public react to a social event? Do negative events have more impact on people than positive events? To answer these questions we need to find the percentage of people in a community experiencing certain emotions and correlate this with the events.
- Identifying changes of emotions in social events: How does the reaction of people in different kinds of events evolves over time? How long does each emotion stay? Ideally, we want to identify what the typical time-variant models are for happiness and sadness. This model will then facilitate mechanisms for predicting changes of emotion in online streams of tweets.
- Assess the effectiveness of our emotion analysis approach: Does it accurately detect evolution of public emotion in a variety of real-life contexts. We deploy our approach in the wild to predict public emotion during different real-life events, including sad and happy events.

II. Proposed Approach to Detect Public Emotion During Events

This section describes our approach to detect public emotion during real-life events. Our approach includes an offline training task and an online classification task. We first develop a system called Emotex to create models for classifying emotion. Then, the created models are deployed in a two-stage framework called EmotexStream to classify live streams of tweets from a specific event.

For the offline task, Emotex is designed to automatically classify each text message into an emotional state. Emotex collects a large dataset of emotion-labeled messages from Twitter. The messages are then preprocessed and converted to feature vectors to train emotion classification models. It then classifies unlabeled messages using the learned models. Our classification algorithms will receive a sample of training points from our labeled dataset $\mathcal{D}$, which we will denote by

$$\mathcal{D} = (t_1, e_1), ..., (t_n, e_n), \quad t_i \in \mathcal{T}, \quad e_i \in E_{class} \tag{1}$$

where $\mathcal{T}$ is the set of all tweets in the labeled dataset $\mathcal{D}$, and E_{class} is the set of emotion labels. Based on the Circumplex model of emotion [11] we defined E_{class} as below:

$E_{class} = \{happy\_active, happy\_inactive, unhappy\_active, unhappy\_inactive\}$

Our emotion classifier is a function that maps a sample tweet t from our test dataset to an emotion class e.

$$e = f(t), \quad t \in \mathcal{T}, \quad e \in E_{class} \tag{2}$$

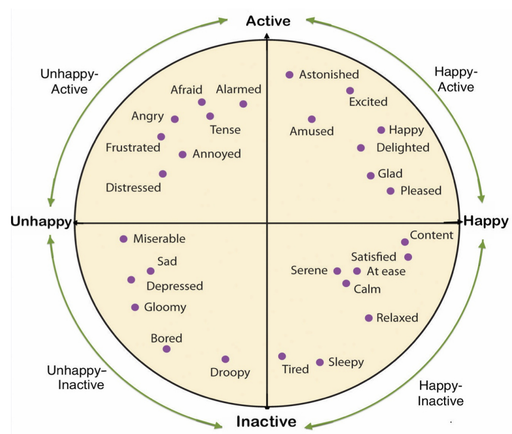

Fig. 1: Circumplex model of affect including 28 affect words by J. A. Russell, 1980. [11]

In the Circumplex model [11], emotions are distributed in a two-dimensional circular space, containing pleasure and activation, as shown in Figure 1. The activation dimension measures if a person is likely to take an action. The pleasure dimension measures how positive or negative a person feels.

Each tweet is first converted into a vector of features. Feature selection plays an important part in the effectiveness of the classification process. We need to capture a set of discriminative and informative features that describe the emotion expressed by each tweet. In our Emotex system, we explored the usage of different features [3]. We use single words, also known as unigrams as the baseline features for comparison.

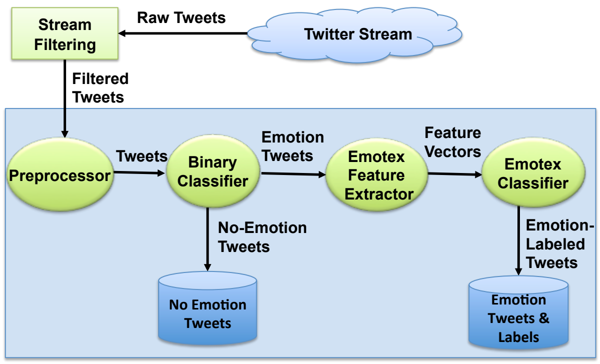

Fig. 2: EmotexStream: A two-stage approach of Emotex for classifying stream of tweets

Emotex utilizes emotion lexicons, as the set of unigram features. Beyond unigrams we also use emoticons, punctuations, and negations as features for emotion classification.

Another key step in emotion classification is the availability of a large set of emotion-labeled data to train classifiers. Manually labeling of Twitter messages with the emotions they express faces numerous challenges, including the inconsistency of human labelers [4]. Therefore, instead in our prior work we have investigated the use of hashtags in Twitter messages that indicate the emotion expressed by tweets as viable alternative to manual labeling [4]. This approach overcomes the need for manual labeling and yields a completely automatic scheme for labeling a massive repository of Twitter messages.

After creating the emotion classifiers, we now deploy the trained models to study public emotion in the wild driven by external social events. However analyzing the real time text is challenging, due to the noise and diversity of tweets in the wild. Since our focus is on emotion detection, we are only interested in processing messages that contain emotions. Thus, we first recognize the tweets representing emotion from the noisy tweets or the ones without emotion. For this purpose, we develop a two-stage framework called EmotexStream for classifying streams of tweets in the wild, fast-paced and voluminous setting.

Figure 2 depicts the two-phase approach of our EmotexStream technology in classifying a general stream of tweets in the wild. As it shows, after cleaning and preprocessing of tweets the first stage categorizes tweets into two general classes, namely *emotion-present* tweets or *emotion-absent* tweets. The second stage utilizes our emotion classification model for a fine-grained emotion classification of tweets with explicit emotion. For binary classification of tweets into either the emotion-present or emotion-absent class we develop an unsupervised method by utilizing emotion lexicons. Our binary classifier assumes that tweets with no emotion are the ones without any emotional or affective words. Therefore, it classifies tweets containing at least one affective or emotional word as emotion-present tweets, while tweets without any affective word are classified as emotion-absent tweets. Different emotion lexicons are available, including ANEW lexicon (Affective Norms for English Words) [12], LIWC dictionary (Linguistic Inquiry and Word Count) [13], and AFINN [14]. We utilize the affective words from these three lexicons and create a comprehensive affective lexicon for our binary classification task.

After binary classification of tweets, emotion tweets then go through the feature selection and multi-class emotion classifier generated by our Emotex technology to classify them based on our defined classes of emotion. We learned the classification model using Support Vector Machine (SVM) algorithm. SVM classifiers partition the data space using linear or non-linear boundaries between different classes. SVMs achieve high performance in text categorization since they accept high-dimensional feature spaces and sparse feature vectors. Also text classification using SVMs tends to be robust to outliers. We used the SVM-light [15] software with a linear kernel to train the SVM classifier.

SVMs are inherently two-class classifiers. The traditional way to do multiclass classification with SVMs is to use one-versus-all or one-versus-one methods to build a set of binary classifiers, and choose the class that is selected by the most classifiers. However, these are not very elegant approaches for solving multiclass problems. Instead, we use $SVM^{multiclass}$ [16]. For linear kernels, $SVM^{multiclass}$ is very fast and its runtime scales linearly with the number of training examples. For a training set $(x_1, y_1), \cdots, (x_n, y_n)$ with labels $y_i \in \{1, \cdots, k\}$, it solves the following optimization problem.

$$\min_{w_1,\cdots,w_k} \frac{1}{2}\sum_{i=1}^{k} w_i^T w_i + \frac{C}{n}\sum_{i=1}^{n}\xi_i \tag{3}$$

subject to:

$$[w_{y1}^T x_1] - [w_y^T x_1] \geq \Delta(y_1, y) - \xi_1, \forall y \in \{1, \cdots, k\} \tag{4}$$

...

$$[w_{yn}^T x_n] - [w_y^T x_n] \geq \Delta(y_n, y) - \xi_n, \forall y \in \{1, \cdots, k\} \tag{5}$$

C is the regularization parameter that trades off margin size and training error. $\Delta(y_n, y)$ is the loss function that returns 0 if y_n equals y, and 1 otherwise.

III. Measuring Public Emotion to Predict Important Moments During the Real-life Events

Twitter messages may refer to a variety of events including, sport games, political elections or natural disasters. These messages may reflect the emotion of users regarding the events. Detecting and measuring public emotion in social networks such as Twitter enable us to identify the influence of social events on a large group of people in near real-time. In this section we describe how to measure public emotion and detect changes in emotion from tweets posted during real life events. The idea is to explore temporal distributions of aggregate emotion during events and detect temporal bursts in public emotion.

Prior research has emphasized the usefulness of social networks and especially Twitter for detecting events [7], [17], [18]. Newsworthy events or sub-events in real life are often captured by detecting temporal bursts of similar messages coming from the same neighborhoods in a social stream [5], [8], [9]. Instead of simply detecting the frequency of similar messages, we are looking for the percentage of people in a community experiencing certain emotions and correlate this with the current event. In fact, our goal is to find temporal bursts of public emotion during real life events. We believe that the temporal burst of the emotional content of tweets can point towards important moments and predict topics of on-going discussion and interest.

In this section, we first introduce the model and definitions for emotion analysis in social streams. Such social streams consist of content-based interactions between users. Then, we describe our online method for measuring public emotion and predicting important moments during the real-life events.

Definition 1 (Tweet Stream): A tweet stream S is a continuous sequence of time-ordered tweet messages $T_1, T_2, \cdots T_r, \cdots$ such that each message contains a text content and is associated with a time and location.

Now, we define an "emotion stream" in the context of tweet stream messages as below:

Definition 2 (Emotion Stream): An emotion stream S_E is a continuous sequence of time-ordered messages $M_1, M_2, \cdots M_r, \cdots$ from a tweet stream, such that each message M_i belongs to a specific emotion class $E_{c1} \in E_{Class}$ (E_{Class} is the set of predefined emotion classes defined in Section II).

In this model, a tweet message M_i in an "emotion stream" can be represented by the tuple $< U_i, T_i, L_i, C_i, E_i >$. This means that a tweet with the message content C_i, has been posted from the location L_i at the given time T_i by the user U_i, and belongs to the emotion class E_i. An example of a tweet belonging to the happy-active class is:

$M_i =< 14,\ Sun\ 2015-02-01,\ 23:59:38,\ [125.33, 65.25],$ "$Well\ done\ Carnival$", $happy\_active >$

The emotion classifier developed by Emotex can be utilized to assign each tweet into an emotion class. Thus, to create an "emotion stream" we apply the classifier model created by our Emotex system to convert a "tweet stream" S into an "emotion stream" S_E. Such an emotion stream typically contains rich information about the public emotion trends.

In order to estimate the value of a specific emotion class E_{c1} among the people in a geographic location L during a time period $< T_1, T_2 >$, we define a function as below:

$$E_{public}(T_1, T_2, L, E_{c1}) = \sum_{T_1 < T_i < T_2} F(M_i, E_{c1}) \qquad (6)$$

where $M_i =< U_i, T_i, L_i, C_i, E_i >$ is a tweet message in the emotion stream with $L_i \in L$, $T_1 < T_i < T_2$, $E_{c1} \in E_{Class}$, and $F(M_i, E_{c1})$ is an indicator function defined as below:

$$F(M_i, E_{c1}) = \begin{cases} 1 & \text{if } M_i \in E_{c1}, \\ 0 & \text{O}therwise. \end{cases} \qquad (7)$$

The occurrence of a real life event may affect the public emotion. During real life events, we can analyze such emotion streams to detect temporal bursts of public emotion. These sudden bursts are characterized by a change in the fractional presence of messages in particular emotion classes. Formally, we define such changes as "emotion bursts", which can point towards important moments during events.

In order to detect emotion bursts, we determine the higher or the lower rate at which messages have arrived to an emotion class in the current time window of length W. The parameters α and β are used to measure this evolution rate.

Definition 3 (Emotion Burst): An emotion burst over a temporal window of length W at the current time T_c is said to have occurred in a geographic region L, if the presence of a specific class emotion E_{c1} during a time period $(T_c - W, T_c)$ is less than the lower threshold α or greater than the upper threshold β.
In other words, we should have either

$$E_{public}(T_c - W, T_c, L, E_{c1}) \leq \alpha \qquad (8)$$

or

$$E_{public}(T_c - W, T_c, L, E_{c1}) \geq \beta. \qquad (9)$$

Now we need to define the upper bound α and lower bound β of public emotion for each emotion class during a temporal window. If our algorithm is applied offline (i.e. tweets from the entire event are available), the thresholds for the entire event can be estimated from the average sum over the whole period of event. However in the online approach the tweets from the entire event are not available. Therefore, in the online approach, we compute the thresholds from the tweets in a temporal sliding window, where the size of the moving window is a parameter.

Let $e_1, \cdots e_i, \cdots e_n$ denote the emotion class E_{c1} of the tweets posted within a temporal window of length W in the emotion stream (n is the number of tweets posted within W). Apparently, $e_1, \cdots e_i, \cdots e_n$ are independent 0-1 random variables (e_i=0 means tweet message M_i doesn't belong to the emotion class E_{c1}, and e_i=1 means tweet message M_i belongs to the emotion class E_{c1}). Based on Equation 6, public emotion within the temporal window W is defined as below:

$$E_{public}(T_c - W, T_c, L, E_{c1}) = \sum_{i=1 \ldots n} F(M_i, E_{c1}) \qquad (10)$$

where $F(M_i, E_{c1})$ is an indicator function of E_{c1}, T_c is the current time and n is the number of tweets posted within window W.

As we know Hoeffding's inequality provides an upper bound on the probability that the sum of random variables deviates $\lambda > 0$ from its expected value as shown by Equation 11:

$$Pr[|X - \mu| >= \lambda] <= 2e^{-2\lambda^2/n} \qquad (11)$$

where X is the sum of independent random variables $X_1, X_2, \cdots, X_n$, with $E[X_i] = p_i$, and the expected value $E[X] = \sum_{i=1...n} p_i = \mu$.

According to the Central Limit Theorem, if n is large then X approaches a normal distribution. We can use Hoeffding's inequality to define an upper bound on the probability that the public emotion E_{c1} deviates from its expected value. Using the Hoeffding bound, for any $\lambda > 0$ we have:

$$Pr[|E_{public}(T_c - W, T_c, L, E_{c1}) - \mu_e| >= \lambda] <= 2e^{-2\lambda^2/n} \tag{12}$$

where μ_e is the expected number of tweets belong to the emotion class E_{c1} in window W and n is the number of tweets posted within W. Given that n is large in a Tweet Stream, emotion class E_{c1} can be approximated using a normal distribution.

$$\mu_e = n \times P_e$$

where P_e is the expected rate of the emotion class E_{c1}.

We use the historical average rate of each emotion class as expected rate P_e for that emotion class. For example, a weekly window can be used to average the rate of each emotion class based on all tweets in general. Therefore, other than a sliding detection window over the recent tweets posted about the event, we also utilize a larger reference window to summarize the past information about the tweets posted in general. In fact, our emotion-burst detection methodology utilizes two sliding windows. One small window W_{event} that keeps the rate of each emotion class based on the most recent tweets posted about the event. Another large reference window $W_{general}$ that keeps the average rate of each emotion class based on all the past tweets posted in general.

Now we describe our methodology to automatically discover emotion bursts during a real life event. First, we create an emotion stream by applying the model created by Emotex system to classify tweets arriving in a stream based on a predefined set of emotion classes. As a second step, our emotion burst detection algorithm then aggregates the tweets of each emotion class into a time-based histogram, using the function in Equation 6. This aggregation allows us to count the rate of each emotion class in each time period.

We then define a sliding window W_{event} (e.g., daily) over the stream of tweets about the event aggregated in temporal bins. We also define a large (e.g., weekly) window $W_{general}$ over the general stream of tweets to keep track of the average rate of each emotion class. In order to perform the burst detection, we continuously monitor the rate of public emotion for each emotion class within each temporal window W_{event}. Whenever the rate of an emotion class exceeds the upper threshold β or falls beneath the lower limit α, an emotion burst is marked as an important moment by keeping its time of occurrence and if it is an up or down case. Then the system signals the occurrence of the detected moments.

IV. Experimental Evaluation: Predicting Important Moments During Real-life Events

In this section we describe the results of measuring public emotion and predicting important moments during different real-life events. In order to study public emotion, we examined emotion distributions by focusing on three emotional contexts; negative, positive and neutral.

A. Data Collection for Event Analysis

Here we describe how we collect Twitter data about several real-life events containing certain emotions, namely sad, happy and angry emotions.

The tweets about an event can be collected by specifying a Twitter keyword query. For example, tweets for a soccer game can be obtained by searching the keywords soccer, football, and team names like Manchester [5].

We analyze two different kinds of real-world events, negative events as well as positive ones. We select the death of Eric Garner in New York [1] as a negative event containing mostly sad and angry emotions. Eric Garner died after a police officer put him in a choke-hold, which caused many discussions on social media. On December 3, 2014, a grand jury decided not to indict the police officer.

We also consider the shooting of Michael Brown on August 9, 2014, in Ferguson, Missouri. Michael Brown was shot by a Ferguson police officer. The disputed circumstances of the shooting and the resultant protests and the civil chaos received considerable attention in the U.S. Following the grand jury announcement, protests, erupted in Ferguson and other cities across the United States in December 2014.

Both of these events stirred public protests and rallies with charges of police brutality. As of December 2014, many demonstrations had been held nationwide against general police brutality. For our analysis purpose, we collected the tweets of users from November 2014 until January 2015 in Massachusetts, New York and Missouri.

As positive events containing mostly happy emotion, we select the New Year 2015 and the Super Bowl game of the National Football League in 2015. For the New Year event we collected tweets from December 22 until January 6. The Super Bowl game was played on February 1, and we collected data from January 26 until February 10.

In order to get information about the emotion of people in general irrespective of particular events, we also collect tweets without any specific hashtags or keywords. In summary, we collect three categories of tweets as described below:

- General tweets: the first category includes general tweets without any specific keyword or hashtag.
- Sad tweets: The second category includes tweets containing hashtags about selected sad events.
- Happy tweets: The third category includes tweets containing hashtags about the selected happy events.

We utilize the Twitter search API to search for a specified set of tweets. Using the Twitter search API we collected tweets

[1] https://en.wikipedia.org/wiki/Death_of_Eric_Garner

Type of Event	Start Date	End Date	#Tweets
Sad Event	Nov 24, 2015	Jan 5, 2015	86K
Happy Event, New Year	Dec 22, 2014	Jan 6, 2015	45K
Happy Event, Super Bowl	Jan 26, 2015	Feb 10, 2015	48K
Neutral	Nov 24, 2015	Jan 5, 2015	67K

TABLE I: Details of the collected tweets

containing one of the hashtags “New Year”, “Christmas”, “XMAS” or ”SuperBowl”. We also collected tweets containing the keywords “Michael Brown” or “Eric Garner”. Table I illustrates our collected tweets in each of three categories.

B. Evaluation Results: Public Emotion During Real-life Events

After collecting different groups of tweets about real-life events, we classify them using our binary classifier to separate tweets with explicit emotion from tweets without emotion (see Section II). Emotion-present tweets will then go through the feature selection and multi-class emotion classifier generated by our Emotex model (see Section II) to classify them into emotion classes. Emotion classes defined by the Emotex system include happy-active, happy-inactive, unhappy-active, and unhappy-inactive (See Section II). In this study, we consider three emotion classes, namely happy (i.e., happy-active), angry (i.e., unhappy-active), and sad (i.e., unhappy-inactive). Then, the emotion-classified tweets are aggregated into a daily-based histogram. Finally, using the methodology described in Section III we are ready to analyze public emotion and detect emotion-critical moments.

The objective of this experiment is to observe the temporal distribution of public emotion during sad and happy events.

Figure 3, presents the changes of different emotion classes during Christmas and New Year 2015 in United States. Changes of public emotion during Super Bowl game are presented in Figure 4. As they show happy emotion during New Year and Super Bowl game is primarily higher than sad and angry emotions by 60% and 40% respectively. During the New Year event the highest rate of happiness is on December 24 and December 31. One day before the Super Bowl play the happiness reaches to its highest rate. These results verifies the effectiveness of our emotion detection approach.

Figures 5 and 6 present the temporal changes of different classes of emotion in New York and Missouri during the selected sad events. The important moments are also specified in each figure. These distributions show a predominance of sad and angry emotions over happy emotion in many days during the sad events. Furthermore, we observe that the distribution of public emotion is relatively similar in different geographic regions.

In order to predict the important moments as emotion bursts, we apply a sliding window W_{event} of length one day over the emotion stream of tweets aggregated in daily bins, as described in Section III. Also a reference weekly window $W_{general}$ is applied over the general stream of tweets to calculate the

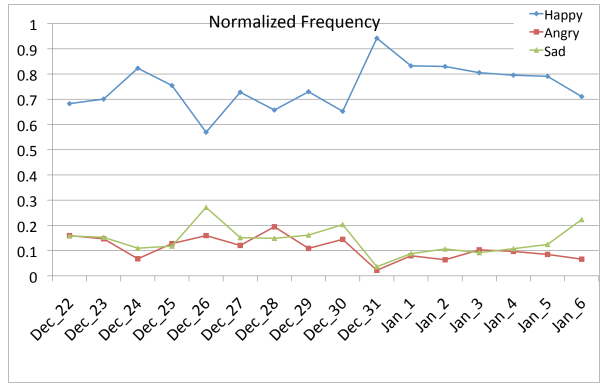

Fig. 3: Changes of emotions during New Year 2015 in United States

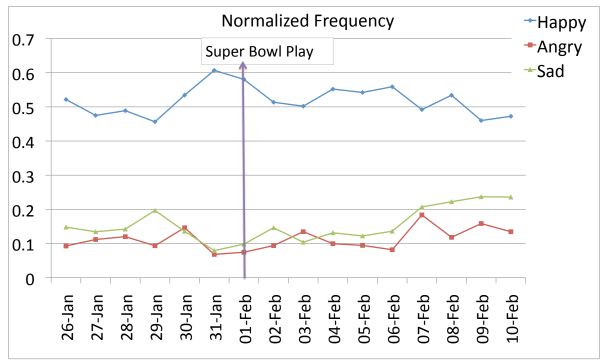

Fig. 4: Changes of emotions during Super Bowl game in United States

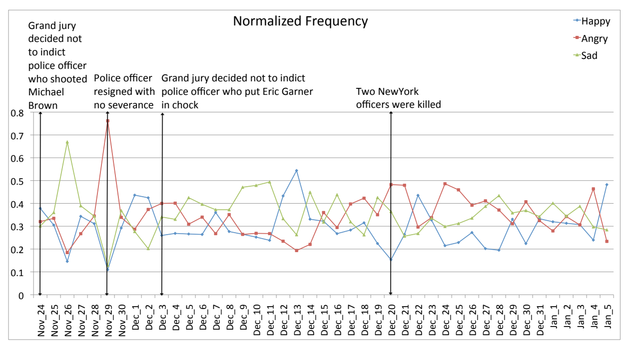

Fig. 5: Changes of emotions about selected sad events in New York

Date	Nov 26th	Nov 29th	Dec 19th	Dec 20th	Dec 27th	Dec 28th	Dec 30th
Happy Rate	210	175	576	462	463	360	503
Boundary (α, β)	(360, 936)	(400, 1040)	(641, 1668)	(753, 1957)	(573, 1491)	(461, 1199)	(561, 1459)

TABLE II: Detected burst changes in happiness

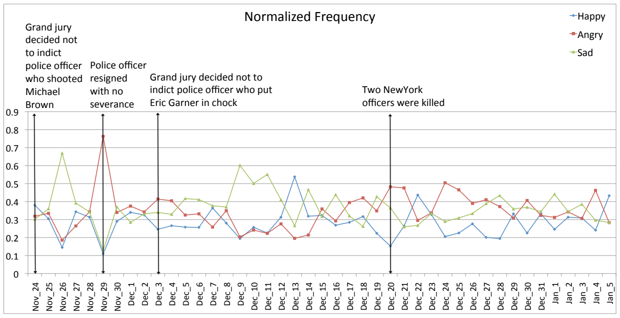

Fig. 6: Changes of emotions about selected sad events in Missouri

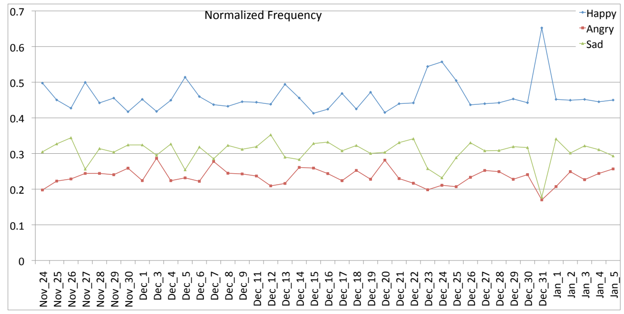

Fig. 7: Changes of general emotions in United States

Date	Nov 29th	Dec 20th	Dec 21st	Dec 24th	Jan 4th
Angry Rate	1219	1453	1420	1211	1247
Boundary (α, β)	(64,704)	(120,1325)	(118,1302)	(99,1096)	(107,1183)

TABLE III: Detected burst changes in angriness

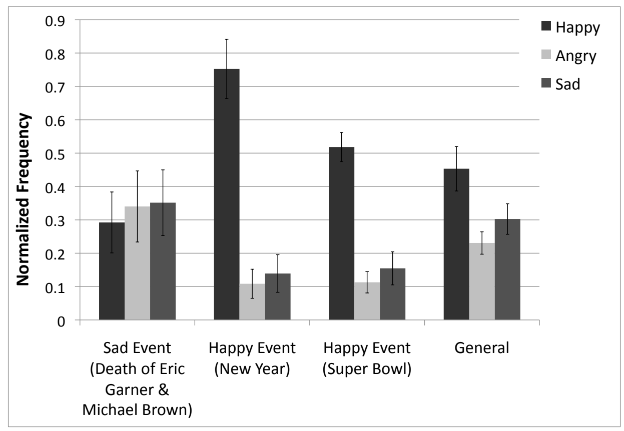

Fig. 8: Distribution of emotions during different events

average rate of each emotion class. Then, we continuously monitor the frequency rate $E_{public}(Tc - W_{event}, Tc, L, E_{c1})$ over time for each emotion class E_{c1}. Whenever this rate for an emotion class exceeds the upper threshold of β or falls beneath the lower limit α, an emotion burst is reported.

Tables II, III, and IV present the days of abrupt changes in happiness, angriness and sadness respectively. The second row shows the frequency rate of emotion bursts which are out of range. The last row shows the low and high boundaries. Comparing the results of these tables with the important moments specified in Figures 5 and 6 confirms that our method is able to detect emotion-critical moments.

Other than sad and happy events, we also collect the tweets in general regardless of particular events. The classification results of the general tweets provide a reference for the average rate of each emotion class. Figure 7 presents the changes of general mood of people in United States. As it is expected the general happiness is higher than the happiness during the sad events, and is lower than the happiness during the happy events.

Table V and Figure 8 show the average percentage and standard deviation of public emotion during different events. As they show, public emotion during happy events fluctuates more than the public emotion in general. However, during sad events the public emotion fluctuates the most (with the standard deviation above 9).

Date	Nov 26th	Dec 11th
Sad Rate	964	659
Boundary (α, β)	(158,734)	(131,651)

TABLE IV: Detected burst changes in sadness

V. Related Work

Recently, there have been substantial research efforts on the identification of events of high importance by monitoring the Twitter stream. Although not directly related, our paper is inspired by many of these prior works in the areas of event detection in Twitter stream.

Prior works have shown the usefulness of real-time analysis of microblogs in many applications via analyzing how the popularity of topics emerges and evolves over time and space [7], [17], [18]. Sakak et al. analyzed microblogs to detect natural disasters. [19]. Kanhabua and Nejdl harnessed microblogging platforms to detect disease outbreaks [20].

Emotion Class	Happy Avg, Std	Angry Avg, Std	Sad Avg, Std
Sad Event	28%, 9.1	35.1%, 10.6	36.2%, 9.8
Happy Event (New Year)	75.2%, 8.8	10.8%, 4.4	13.9%, 5.6
Happy Event (Super Bowl)	51.8%, 4.4	11.3%, 3.2	15.5, 5
General	45.3%, 6.6	23.1%, 3.3	30.2%, 4.5

TABLE V: Average and standard deviation of emotion classes during different events

A lot of works on detecting important moments or events are based on the volume of similar tweets [5]–[9]. Marcus et al. developed TwitInfo for visualizing and summarizing events on Twitter [5]. TwitInfo automatically detects temporal peaks in tweet frequency, add labels to the peaks, and visualize them in a timeline. Also, it displays sentiments of tweets in a pie chart.

Bifet et al. [21] developed a real-time system to read tweets and detect changes by finding the terms whose frequency changed. They use ADWIN algorithm [22] to detect changes. Weng et al. [7] developed an event detection algorithm named EDCoW. They use wavelet analysis on the frequency of the words from tweets. Then they cluster the words with high signal-autocorrelations to form events.

Mathioudakis and Koudas [23] developed a system that detects an event when a set of keywords appear together at an unusual and high rate. Their work relies on offline analysis, which is not suitable for real-time analysis.

Valkanas and Gunopoulos [24] developed a system that clusters a group of users according to their geographical location and then monitors the emotion of each group. Similar to ours, their system reports an event, when the group's cumulative emotion changes abruptly. However, they approximate the Probability Density Function of the aggregate emotion to detect abrupt change of emotion.

Nichols et al. [6] summarized Cup soccer matches and detect sub-events when the volume of status updates exceeds a threshold value. This value is computed offline from basic statistics of the set of all slopes for that match. They also presented an online approach where the threshold is computed using a sliding window.

VI. Conclusion

In this paper we analyze public emotion trends driven by social events and investigate its temporal distributions using massive microblogs on twitter. We propose an online approach to measure public emotion and predict important moments during social events. We deploy our approach in the wild to predict public emotion during different types of real-life events (i.e., pleasant and tragic events). From the daily tweets we were able to observe interesting temporal changes in public positive and negative emotion and also identified major moments when public emotional tweets are intensive.

References

[1] M. De Choudhury, S. Counts, and M. Gamon, "Not all moods are created equal! exploring human emotional states in social media," in *6th International AAAI Conference on Weblogs and Social Media*, 2012.

[2] H. Kwak, C. Lee, H. Park, and S. Moon, "What is twitter, a social network or a news media?" in *Proceedings of the 19th international conference on World wide web*. ACM, 2010, pp. 591–600.

[3] M. Hasan, E. Rundensteiner, and E. Agu, "Emotex: Detecting emotions in twitter messages," in *Proceedings of the 6th ASE International Conference on Social Computing*. Academy of Science and Engineering (ASE), USA, 2014.

[4] M. Hasan, E. Agu, and E. Rundensteiner, "Using hashtags as labels for supervised learning of emotions in twitter messages," in *Proceedings of the ACM SIGKDD Workshop on Healthcare Informatics*, 2014.

[5] A. Marcus, M. S. Bernstein, O. Badar, D. R. Karger, S. Madden, and R. C. Miller, "Twitinfo: aggregating and visualizing microblogs for event exploration," in *Proceedings of the SIGCHI Conference on human Factors in Computing Systems*. ACM, 2011, pp. 227–236.

[6] J. Nichols, J. Mahmud, and C. Drews, "Summarizing sporting events using twitter," in *Proceedings of the 2012 ACM International Conference on Intelligent User Interfaces*. ACM, 2012, pp. 189–198.

[7] J. Weng and B.-S. Lee, "Event detection in twitter," in *PProceedings of the Fifth International Conference on Weblogs and Social Media*. The AAAI Press, 2011, pp. 401–408.

[8] C. Aggarwal and K. Subbian, "Event detection in social streams." in *Proceedings of the SIAM International Conference on Data Mining*. SIAM, 2012, pp. 624–635.

[9] D. A. Shamma, L. Kennedy, and E. F. Churchill, "Peaks and persistence: modeling the shape of microblog conversations," in *Proceedings of the ACM 2011 conference on computer supported cooperative work*. ACM, 2011, pp. 355–358.

[10] F. Chierichetti, J. Kleinberg, R. Kumar, M. Mahdian, and S. Pandey, "Event detection via communication pattern analysis," in *Proceedings of the 8th International AAAI Conference on Weblogs and Social Media*, 2014.

[11] J. A. Russell, "A circumplex model of affect," *Journal of Personality and Social Psychology*, vol. 39, pp. 1161–1178, 1980.

[12] M. M. Bradley and P. J. Lang, "Affective norms for english words (anew): Instruction manual and affective ratings," Citeseer, Tech. Rep., 1999.

[13] J. W. Pennebaker, M. E. Francis, and R. J. Booth, "Linguistic inquiry and word count: Liwc 2001," *Mahway: Lawrence Erlbaum Associates*, vol. 71, 2001.

[14] F. rup Nielsen, "A new anew: evaluation of a word list for sentiment analysis in microblogs," in *Proceedings of the ESWC2011 Workshop on 'Making Sense of Microposts': Big things come in small packages*, 2011, pp. 93–98.

[15] T. Joachims, *Making Large-Scale SVM Learning Practical*. MIT Press, 1999.

[16] ——, "Multi-class support vector machine," https://www.cs.cornell.edu/people/tj/svm_light/svm_multiclass.html, 2008.

[17] 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 (CIKM)*. ACM, 2012, pp. 155–164.

[18] A. Ritter, O. Etzioni, S. Clark *et al.*, "Open domain event extraction from twitter," in *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 2012, pp. 1104–1112.

[19] 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*. ACM, 2010, pp. 851–860.

[20] N. Kanhabua and W. Nejdl, "Understanding the diversity of tweets in the time of outbreaks," in *Proceedings of the 22nd international conference on World Wide Web companion*, 2013, pp. 1335–1342.

[21] A. C. Bifet Figuerol, G. Holmes, B. Pfahringer, and R. Gavaldà Mestre, "Detecting sentiment change in twitter streaming data," *Journal of Machine Learning Research*, pp. 5–11, 2011.

[22] A. Bifet and R. Gavalda, "Learning from time-changing data with adaptive windowing." in *Proceedings of the SIAM International Conference on Data Mining*, 2007.

[23] 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*. ACM, 2010, pp. 1155–1158.

[24] G. Valkanas and D. Gunopulos, "How the live web feels about events," in *Proceedings of the 22nd ACM international conference on Information and Knowledge Management*. ACM, 2013, pp. 639–648.