Sentiment Monitoring of Social Media from Oceania

Ross Stewart Sparks

Received: 15 December 2017 Accepted: 4 January 2018 Published: 15 January 2018

Abstract

Introduction-Social media platforms have experienced a major growth in the past few years, with people choosing to communicate, very often publicly, through social media. They disseminate information, opinions, and announcements. They also share a lot about themselves and their experiences. In particular, they often share information about how they feel. This potentially provides a wealth of information, in real-time, about the emotional state of individuals or communities. This can, in turn, provide valuable information about how people react to various events.In our work, we have been investigating whether we can process emotion-related information from social media in real time, to understand how people react to different events and circumstances and potentially also help further research in mental health. To this end, we developed We Feel, a tool that analyses emotions on Twitter and presents them through an interactive visualization (see wefeel.csiro.au). We Feel constantly monitors the Twitter stream, looking for tweets (in English) containing any emotional content (Paris et al., 2015;Larsen et al., 2015). The platform aims at monitoring the regional elevated risks of suicide by assessing the mood of people in that region.

***Index terms*—**

1 Sentiment Monitoring of Social Media from Oceania

Ross Sparks ? & Cecile Paris ? I. Introduction ocial media platforms have experienced a major growth in the past few years, with people choosing to communicate, very often publicly, through social media. They disseminate information, opinions, and announcements. They also share a lot about themselves and their experiences. In particular, they often share information about how they feel. This potentially provides a wealth of information, in real-time, about the emotional state of individuals or communities. This can, in turn, provide valuable information about how people react to various events.

In our work, we have been investigating whether we can process emotion-related information from social media in real time, to understand how people react to different events and circumstances and potentially also help further research in mental health. To this end, we developed We Feel, a tool that analyses emotions on Twitter and presents them through an interactive visualization (see wefeel.csiro.au). We Feel constantly monitors the Twitter stream, looking for tweets (in English) containing any emotional content (Paris et al., 2015;Larsen et al., 2015). The platform aims at monitoring the regional elevated risks of suicide by assessing the mood of people in that region. Figure 1 shows a screen shot of We Feel. The set of emotions that are captured is shown on the left, displayed as an "emotion wheel". A map of the world is on the right. Both of these elements are interactive: one can select a region in the world, or a specific emotion, and the visualisation in the centre will focus on the chosen attributes (location or emotion) and change accordingly. In Figure 1, a specific date (May 21-27, the week of the Manchester attack), region (Oceania) and emotion (sadness) have been chosen. The visualisation shows the emotions as reflected in the tweets being processed, colour-coded by emotions, matching the wheel.

In this paper we use We Feel to explore the mood of the people in Oceania (Australia and New Zealand) over the period running from1 June 2014 to the end of November 2016. This paper uses statistical process control to flag significant changes in the mood of a region and understand its implication on the society in that region. We are interested in what events influenced the mood.

An event may dominate the public conversation, so the number of people that talk about it increases significantly when it occurs, and then subsides as people either lose interest, all the issues of the event are people's interest, or life simply moves on. The monitoring technology in this paper is interested in isolating the dominant sentiment for an event. An event is determined by a significant increase of the number of tweets. The dominant sentiment for an event is found by monitoring the proportion of tweets with sentiments classified as expressing either anger, fear, surprise, sadness, joy or love. The final aim is to understand when people respond to events, why they respond with certain sentiments and how quickly does the event stop influencing the mood of people, or in other words how quickly do people move on with their lives after an event.

2 II. Event Detection

We start by detecting an event. As mentioned above, an event is defined as an unusual increase in the number of tweets per hour. We thus first need to define what is usual before we can establish what is unusual. We used the total tweets per hour (See Figure 2) as a response variable with explanatory variables lag logarithm hourly counts, time, harmonics to model both seasonal trends and within hour trends, and day-of-the week influences. Public holidays are ignored because the region does not have consistent public holidays. We assumed that the harmonic for season and day interacted. This model fitted quite well with the Pearson residuals showing no significant autocorrelation. The EWMA chart applied to the Pearson residuals of this model looked very strange with it mostly hugging the centreline and with no high-sided signal. Further investigation revealed that the lag 1 autocorrelation in the hourly counts was not very high at 0.54, and the coefficient for the logarithm lag counts in the fitted model was 0.308. This autocorrelation was driven by the events where counts ramped up. However, while they communicated with friends between events, there was no apparent autocorrelation until the next event. For this reason, we decided to fit the above model without the explanatory variable lag logarithm hourly counts included, and used this model to define usual behaviour. This meant that we would live with a slightly higher over-dispersion in the model than is justified, because we have included all events in the model without accounting for their autocorrelation, but we were happy to live with that and only focus on the major events. These total hourly tweets appear to be over dispersed with a number of low and high sided outliers. We are interested in detecting the high sided outliers which we try to associate with a historical event that we believe created the significantly elevated interest amongst Twitter users. To achieve this, we apply the EWMA chart to the Pearson residuals for the model above. Firstly we establish the expected total hourly tweets by fitting the negative binomial regression model defined above. We estimate the Pearson residuals for this model and then apply the EWMA chart with exponential weights given by 0.4 because most events seem to wane very quickly in the social media context, and most of the events we are looking at are fairly large shifts. We believe that this is appropriate because most Twitter users' attention span is fairly short, seemingly less than an hour.

3 Fig. 3: Allocation of High-Sided Signals to an Event

We applied an EWMA control chart to the Pearson residuals to flag the unusual events of the study period using a retrospective surveillance approach. The in-control Average Run Length (ARL) for this EWMA was taken as 365 in designing the plan. The threshold was found by simulation, but we could have used the spc package in R (Knoth, 2017) to provide a very similar threshold. Since we are dealing with hourly data, this gives us roughly 24 false alarms on average per year. Figure 4 provides the results of this chart by signalling unusual events: they occur outside the upper dashed red line either on the high-side or the low-side. We will ignore the low-sided signals in Figure 4 (events that trend below the red dashed line).

4 III. Understanding the Twitter Posts'

Sentiments for the Events Each tweet is classified as having one (or more) of the following sentiments: anger, fear, joy, love, sadness or surprise. We are interested in two cases: (1) when there is a change in sentiment independently of whether there is an event of not; and (2) to explore the sentiments for the events discovered in the previous section. In this section, we explore the first scenario. The second scenario demands a multivariate approach; it will be explored in the next section. Here, we are interested whether the sentiments change significantly over time independently. To carry this out we fit the following model using fear as an example. The modelling is then identical for all other sentiments. Fear: The resulting EWMA chart for fear is included in Figure **??**. This flags three events where fear was significantly higher than usual: (1) (3) an increased proportion of angry tweets on 12 July 2015; and (4), again, on 9 November 2016.

5 Fig. 5: Unusual Proportion of Tweets Expressing Anger

Surprise: Now we explore tweets that express a higher than expected proportion of tweets with sentiment surprise (see Figure **??**). We see 7 peaks of surprises. The first surprise is, I am guessing, during the protests at the G20 summit in Brisbane. The second is when 2 of the Bali 9 drug smugglers jailed in Indonesia where executed by a firing squad. The third was Johnny Depp illegally smuggling his dogs into Australia from the USA.

The forth is Penrith teenager caught with a gun in a school in a western suburb of Sydney. The fifth was Russia starting to attack ISIL in Syria. The sixth is the climate pact agreement, which seems to last a long-time when most other events seem to dissipate quite quickly. The last is a massive shift from low surprise to massive surprise on the BREXIT election outcome.

6 IV. Multivariate Views of the Sentiment Analysis

In order to understand the mood of Australians during the study period, we need a multivariate view of the sentiment monitoring process. The first multivariate view of the sentiment counts is achieved using parallel coordinate plots. An example is displayed in Figure 11. It displays the full list of sentiment counts for 6 days jointly using a parallel coordinate plot. This allows us to jointly view trends for all sentiment counts in a single plot, displaying trend information for all sentiment counts relative to their expected values. The lines go from black being the most recent date (14 November 2016), followed by red, green, blue, light blue, magenta and yellow (9 November 2016). The confidence bounds are the thresholds for the EWMA statistic for the sentiment scores. This plot helps us identify that there is a rough trend regional counts towards greater volumes expressing anger, fear and sadness and a reduction in joy and love. Note that love started with an unusually high number of counts This plot is easy to interpret and helps interpret the full picture of the sentiment scores. It does not, however, make the best use of the relationships between the variables/sentiments. To capture this relationship, we propose using the dynamic biplot of Sparks et al. (1997). It monitors changes in location of the counts as well as changes in correlation between the tweet counts and changes in dispersion of the counts in a single plot, making it quite useful in interpreting the Twitter users' responses to certain events. For example, Figure 11 describes the response to the shooting down of flight MH17 over Ukraine. Note that 85% of the variation is in two dimensional display; 58% in the first dimension and 27% in the second dimension. The overwhelming response is one of sadness and significantly reduced joy. There is a significant increase in fear and anger but this is roughly orthogonal to those that express sadness. Note that many people are expressing anger and fear at the same time, as we see that these two emotions are close to being collinear. There was also a simultaneous reduction in the expression of joy, mostly from those that expressed sadness. The correlation between these sentiments counts have not changed significantly by the colours in the matrix below the variable plot. We conclude that the initial response to Phillip Hughes's death was a mixture of sadness and anger; but, later (on represented graphically), as people like Michael Clarke (the then Australian cricket captain) expressed his mateship for Phillip Hughes, this changed to the dominant response becoming love for the man who had so tragically lost his life. Figure 14 indicates that the dominant response to Rosie Batty becoming Australian of the year was one of surprise, and all other sentiments were orthogonal to this, indicating that no other sentiment increased. This is fascinating, but it is unclear whether people were surprised about Tony Abbott (then Australian Prime Minister) making such a call, or whether they were surprised by the choice of Rosie. This choice did raise the serious issue of domestic violence within Australia, and Rosie was the perfect ambassador fighting against domestic violence seeing she had experienced it firsthand (she, and many others, witnessed her ex-husband killing their son after a cricket match). Note that there was no change in the correlation structure indicated by the matrix of boxes below the variable plot not being coloured.

In Figure 15, the dominant response to the energy debate after the South Australia energy crisis (a total blackout after a major storm) was one of increased sadness, with no other sentiment increased. The issue was one where severe weather-related events cut the supply of energy to the entire state, which has a large proportion of renewable energy. This started a national debate about the state relying too much on renewable energy sources. The interesting feature of this response was that there was no increase in angry tweets because of the state government's decision on the percentage of renewable energy to be used. I think this means that the South Australian residents don't strongly disagree with the South Australian state government energy policy. Note that there was a change in correlation structure with love and joy became less correlated. The other colours indicate warnings.

7 K

In Figure 16, the dominant but weak response to the news that Australian Airforce Jets were starting to operate in Syria for the first time was initially one of anger, but this did not last long; no more than a few hours before the response was an increase in joy was dominant and remain so for more than the next 24 hours. This increase in joy was not massively significant because the mean square error for joy did not flag as significant (the joy line in the vector plot was not coloured red but the sausage shape in the middle of this vector indicate a significant increase in joyful responses). Note that there was a change in the correlation structure: love and joy became less correlated, and love and anger became more positively correlated as these counts both decreased simultaneously.

8 Fig. 15: Twitter response to the Australian Airforce Jet operating in Syria

The Twitter response to the arrest of Gino & Mark Stocco (Father and son) after being on the run for 8 years was a strong response of sadness, and this is mostly driven by two hours of the day at about 6 and 7pm at night when the arrest was probably reported. This does not make a whole lot of sense, but there was a nonsignificant reduction in the surprise, love and joy tweets which makes more sense when harden criminals are arrested. Potentially this was a case of things going wrong for two Aussie battlers.

In Figure 17, the response to phone data retention laws for internet service providers in Australia was one of increased sadness and reduced joy, but the observation plot does not flag a multivariate shift in location. Thus this response is not very strong. There is no change in correlations.

9 Conclusion

We have demonstrated ways of monitoring tweet sentiment scores for a region as a way of understanding how the region responds to events. We first defined events as those periods where the number of tweets for the region significantly increased. We then monitored how unusual the counts of these tweets were after correcting for the volume of tweets. This was achieved for each sentiment independently; however, these sentiment counts are correlated, and monitoring them independently makes interpreting the response to events quite difficult. The parallel coordinate plots are relatively easy to understand. They display trends in a reasonable way but ignore correlations. Therefore we prefer the dynamic biplot which monitors changes in location, dispersion and correlations simultaneously in one plot. It is also efficient at displaying trends in the observation plot. Although its interpretation is complex, we believe the rich information it presents makes it a reasonable tool for monitoring and understanding events. [1]

1
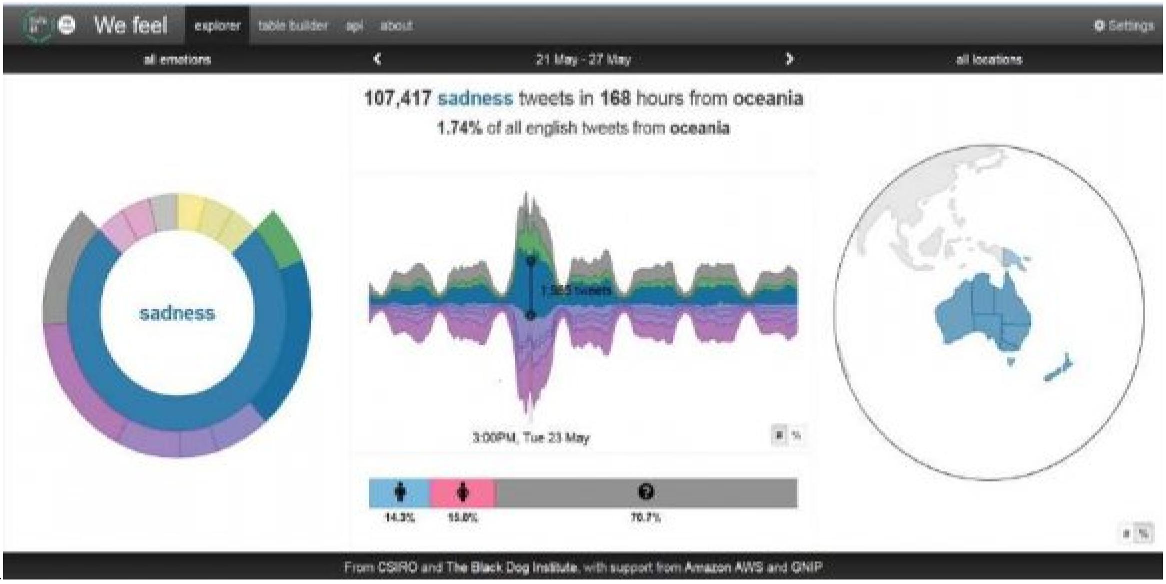

Figure 1: Fig. 1 :

[1] © 2018 Global Journals K

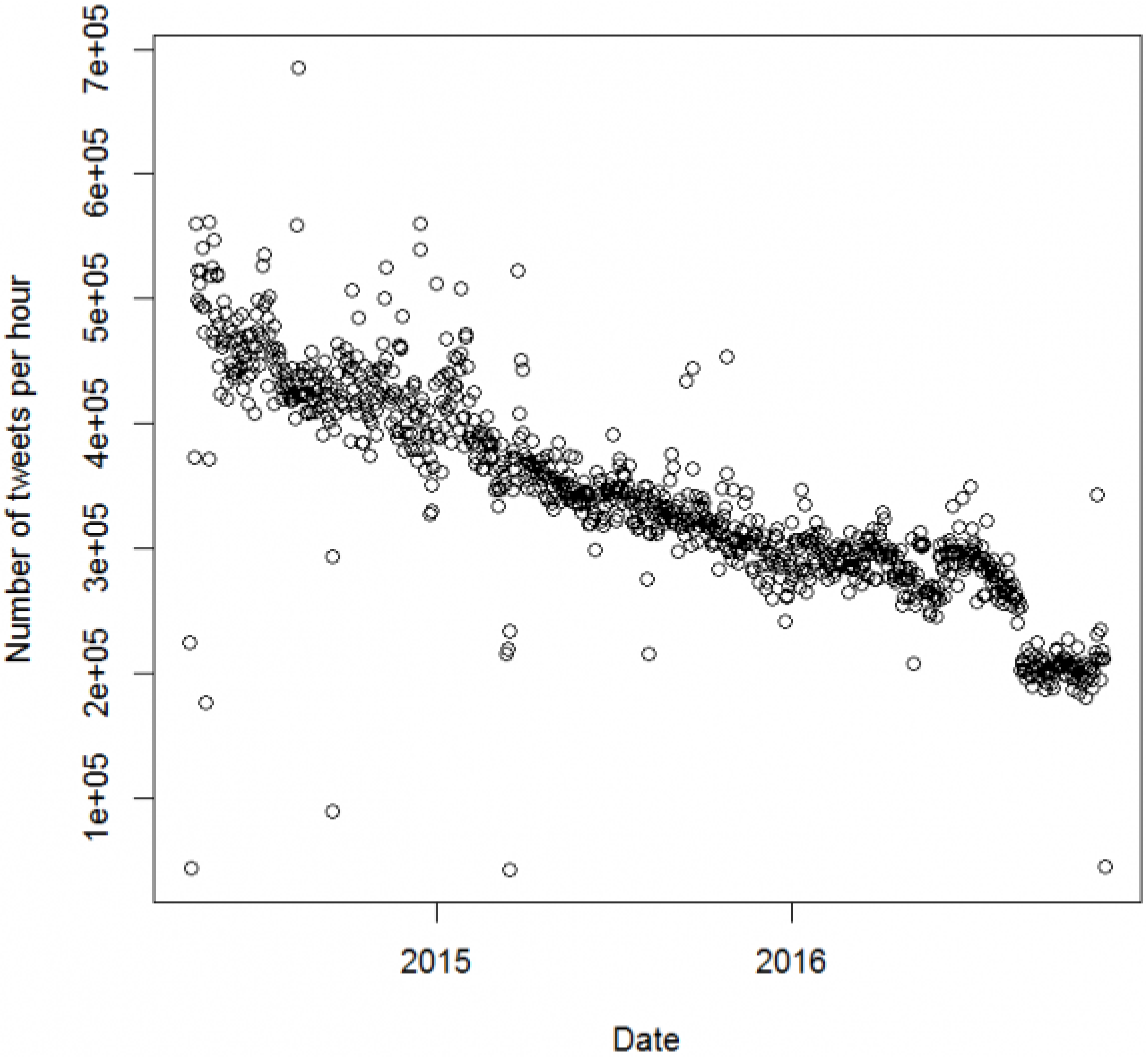

1

Figure 2: Fig. 1 :

3

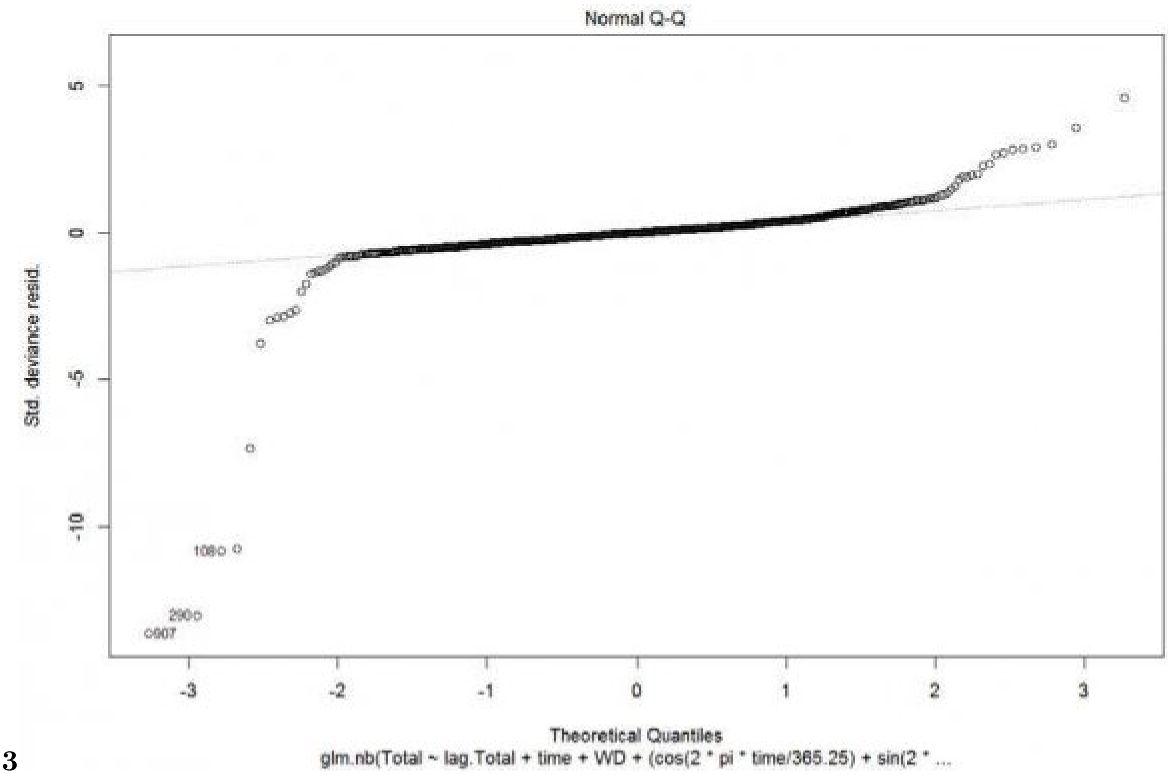

Figure 3: Figure 3

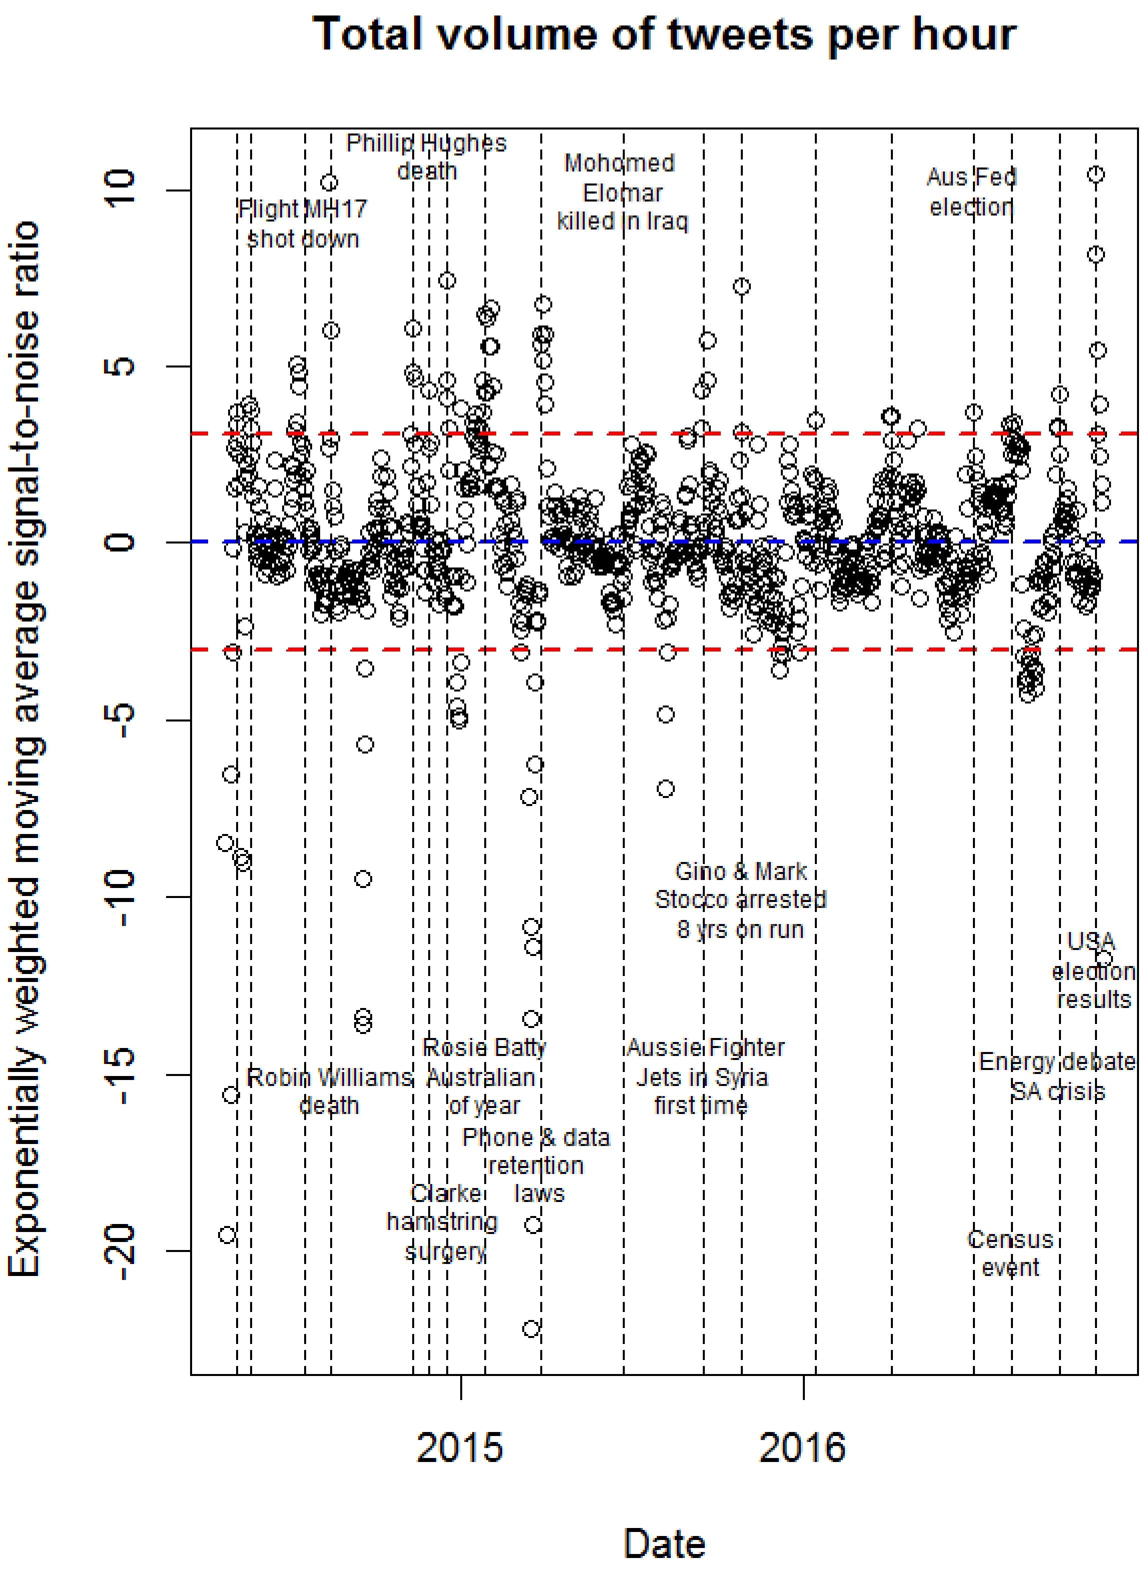

2

Figure 4: Fig. 2 :

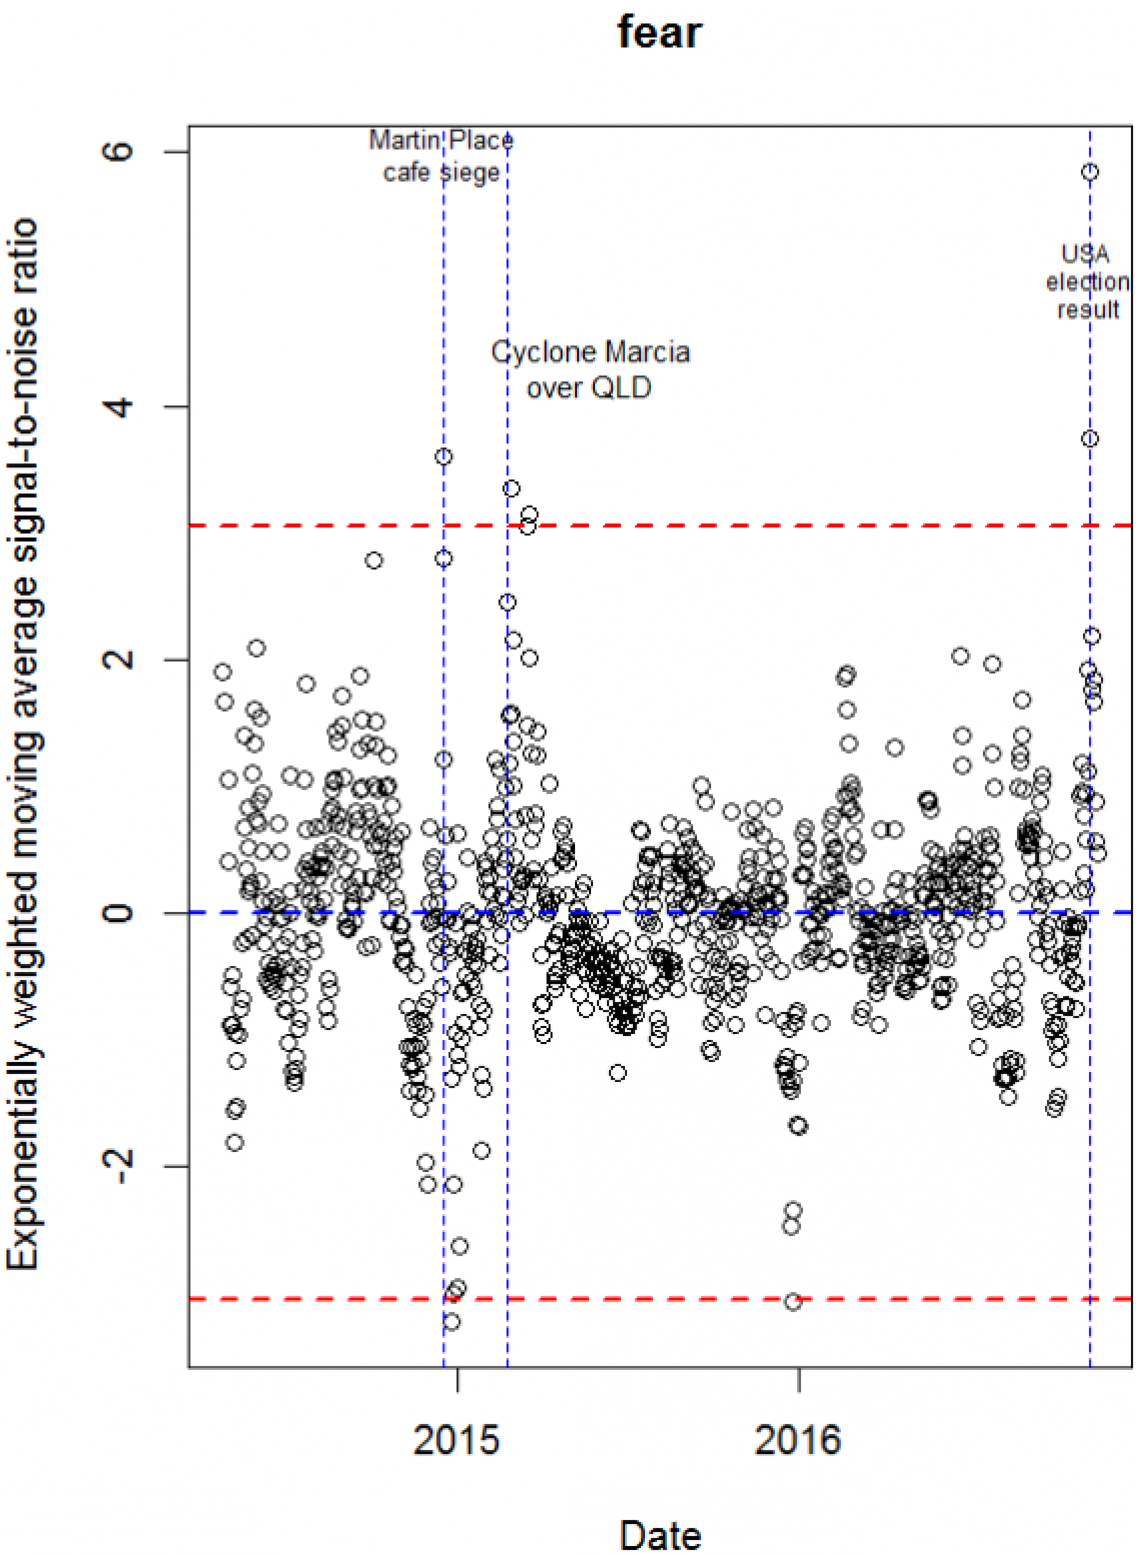

4

Figure 5: Fig 4 :

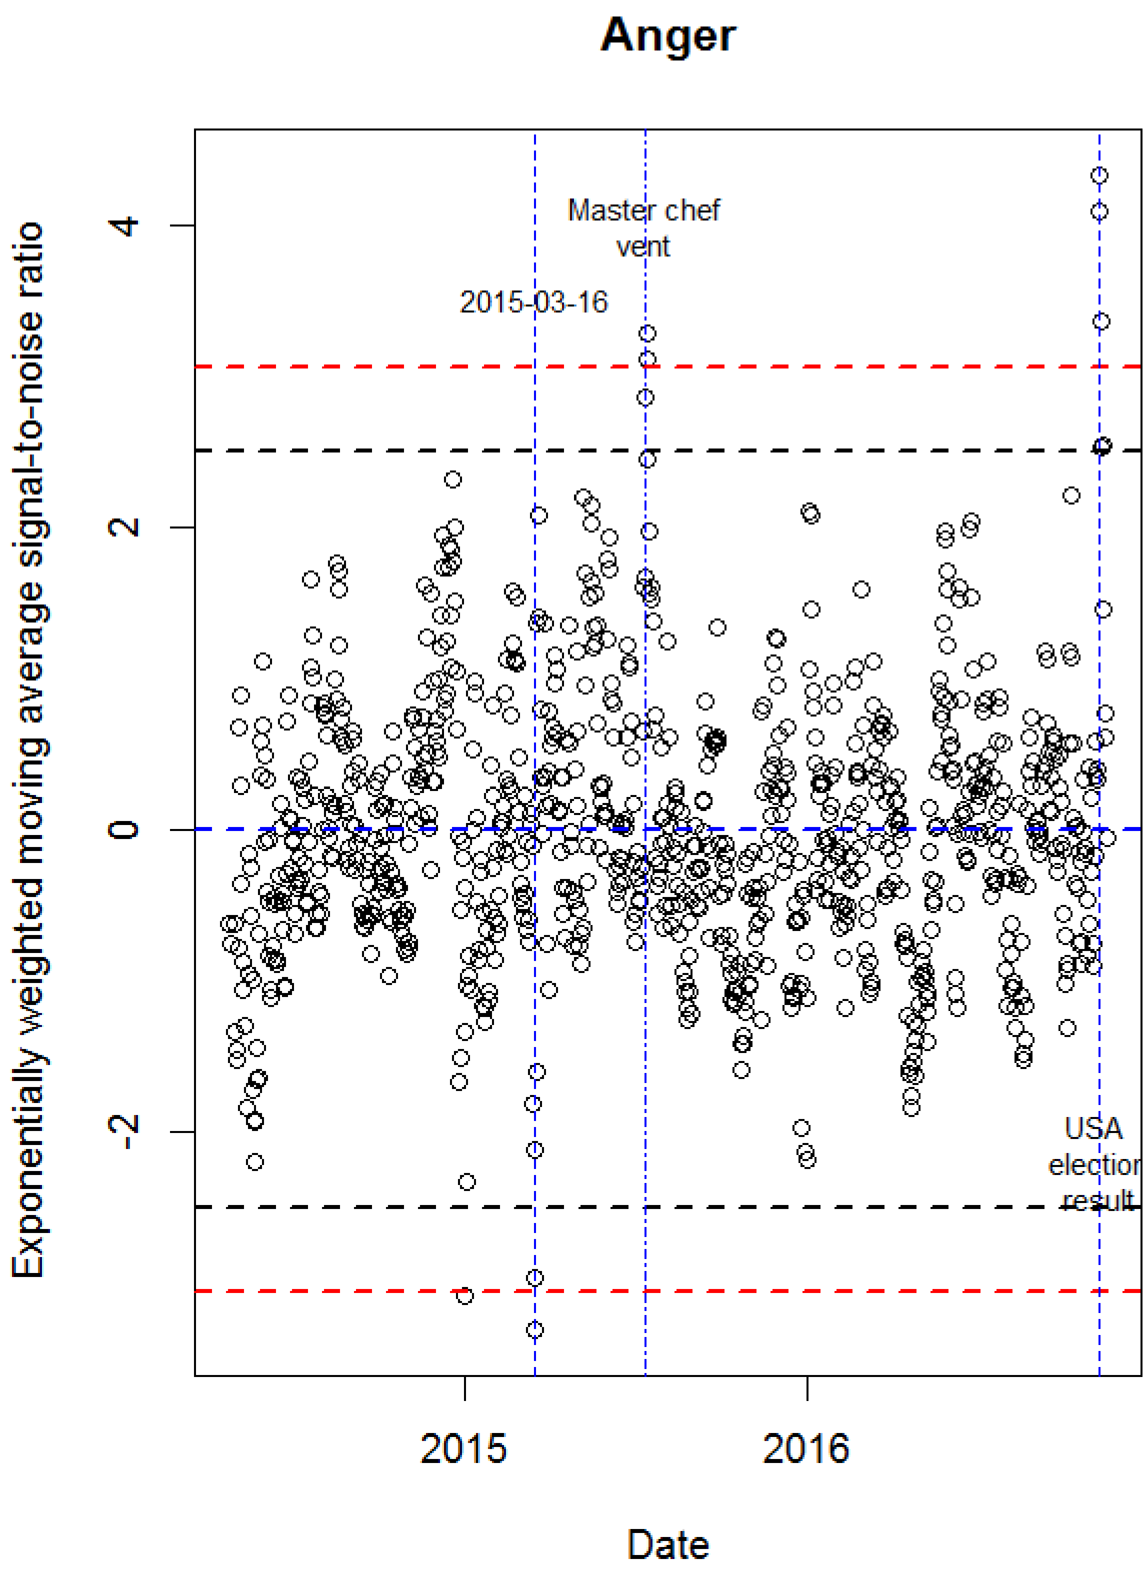

61789

Figure 6: Fig. 6 : 1 KFig. 7 :Fig. 8 :Fig. 9 :

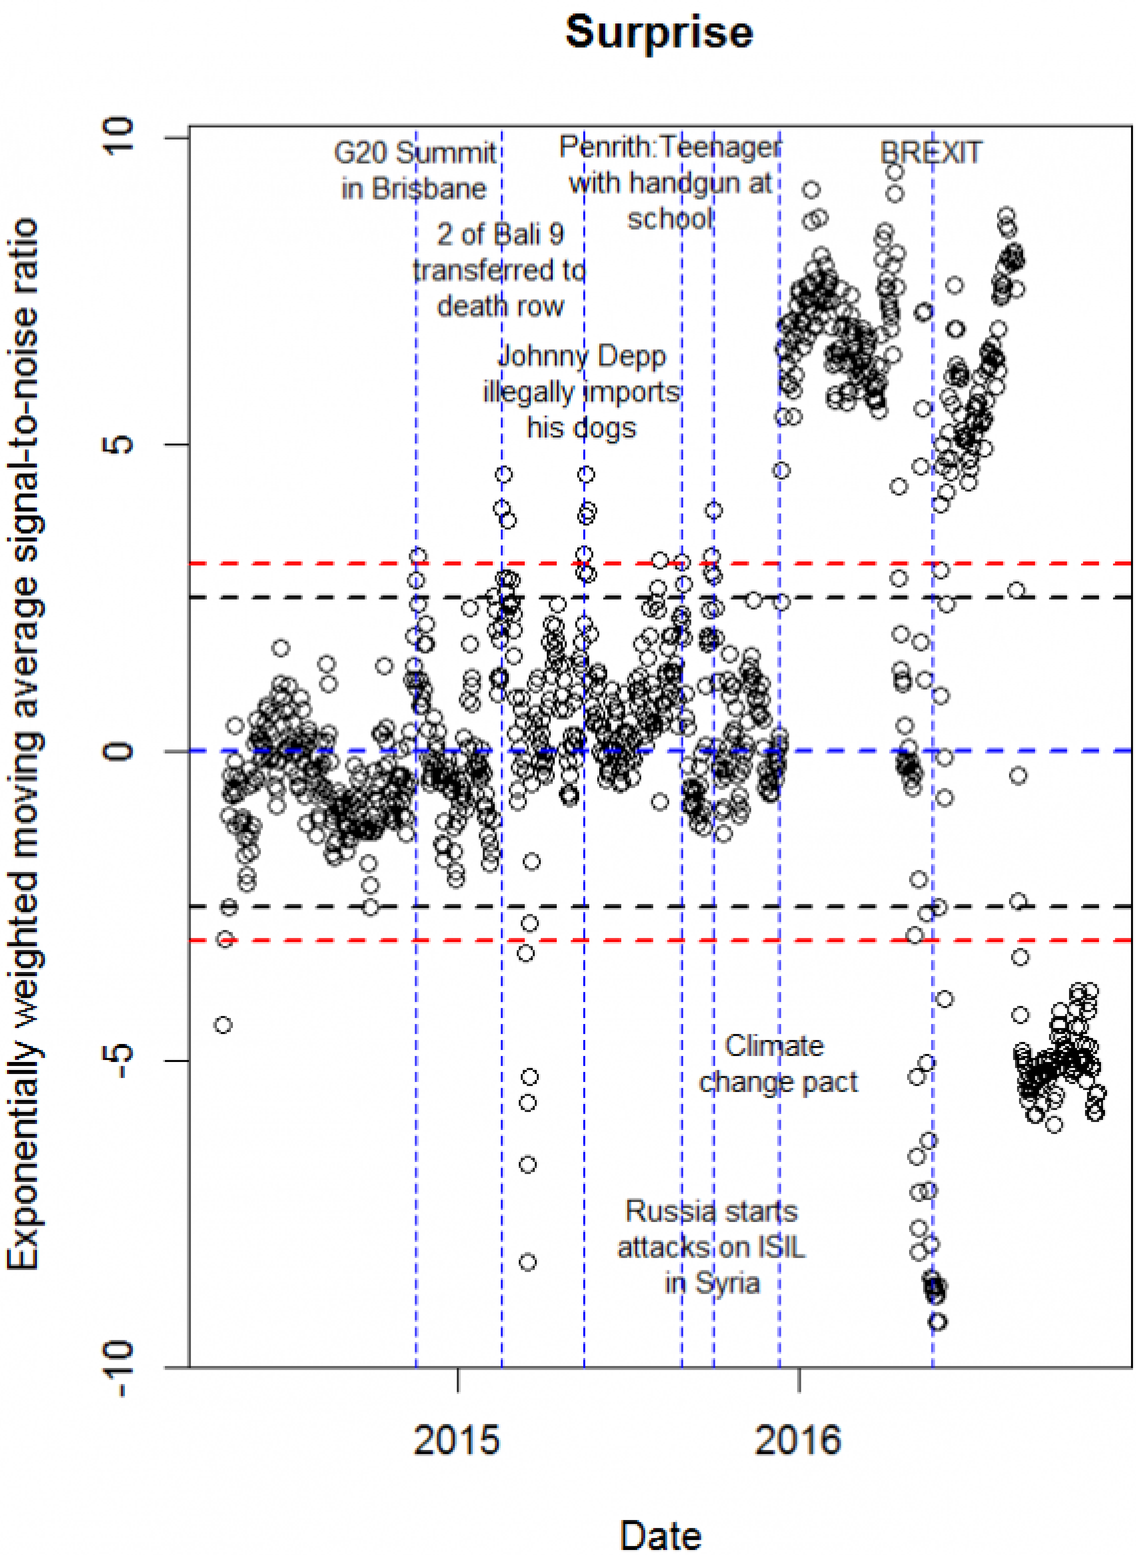

10

Figure 7: Fig. 10 :

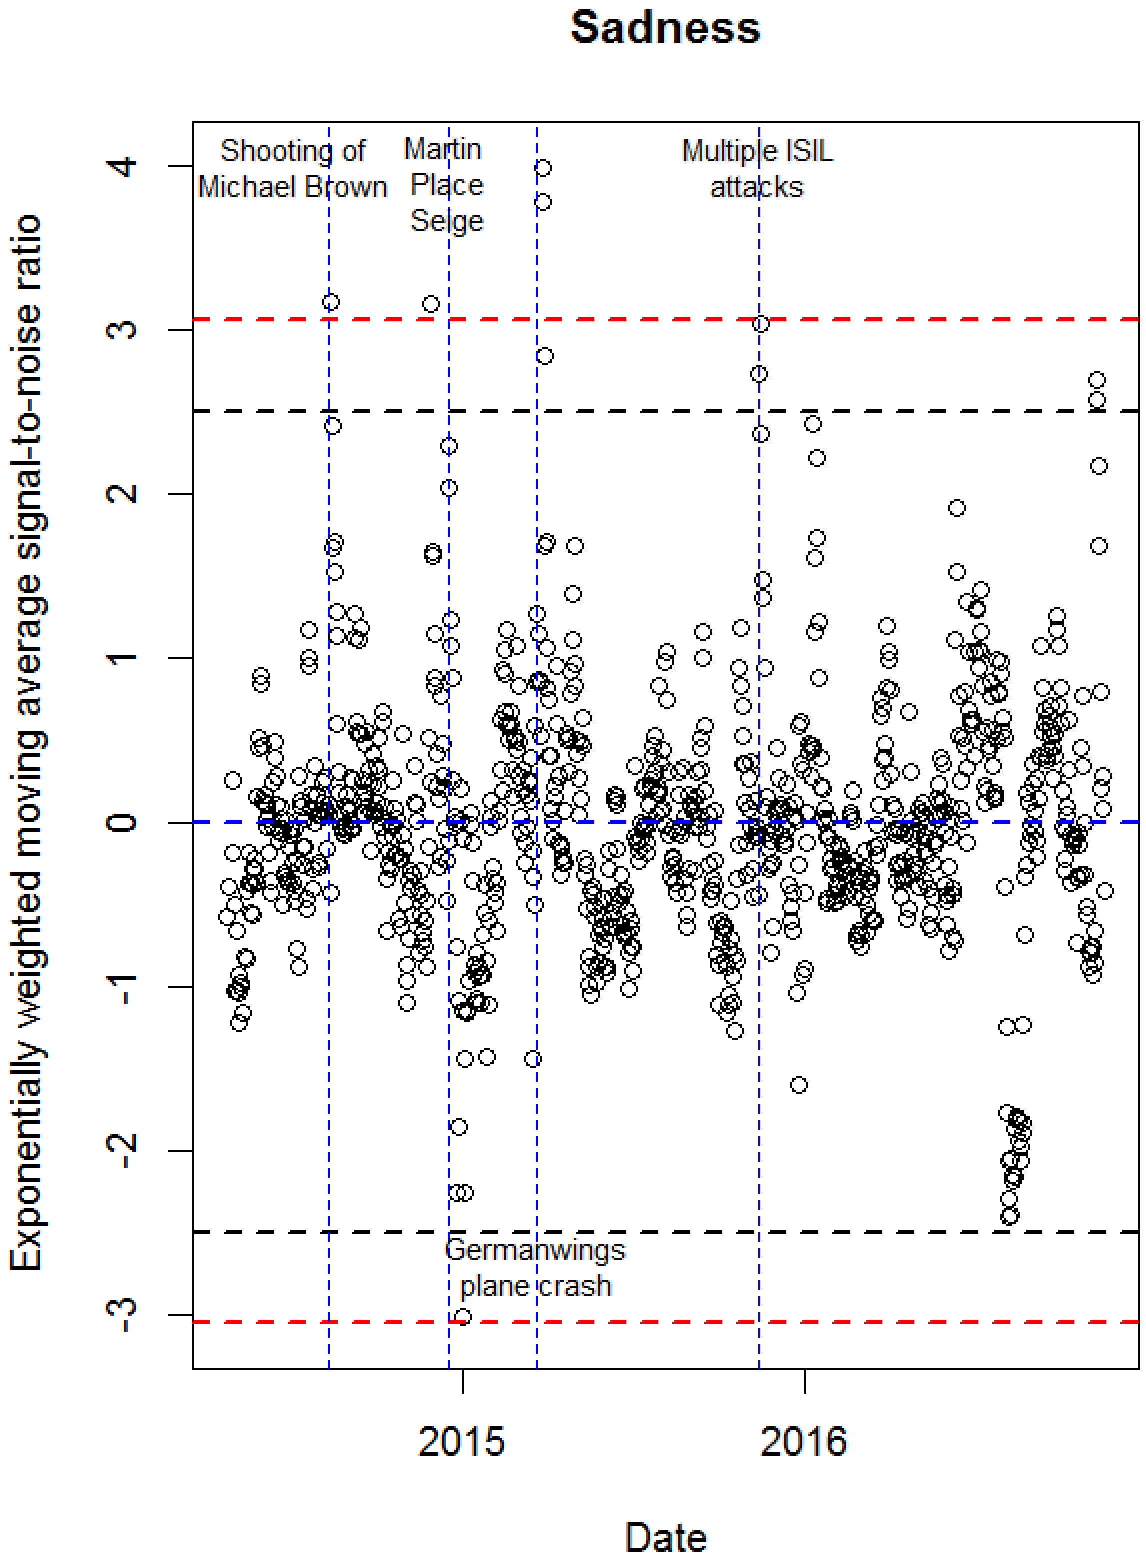

11

Figure 8: Fig. 11 :

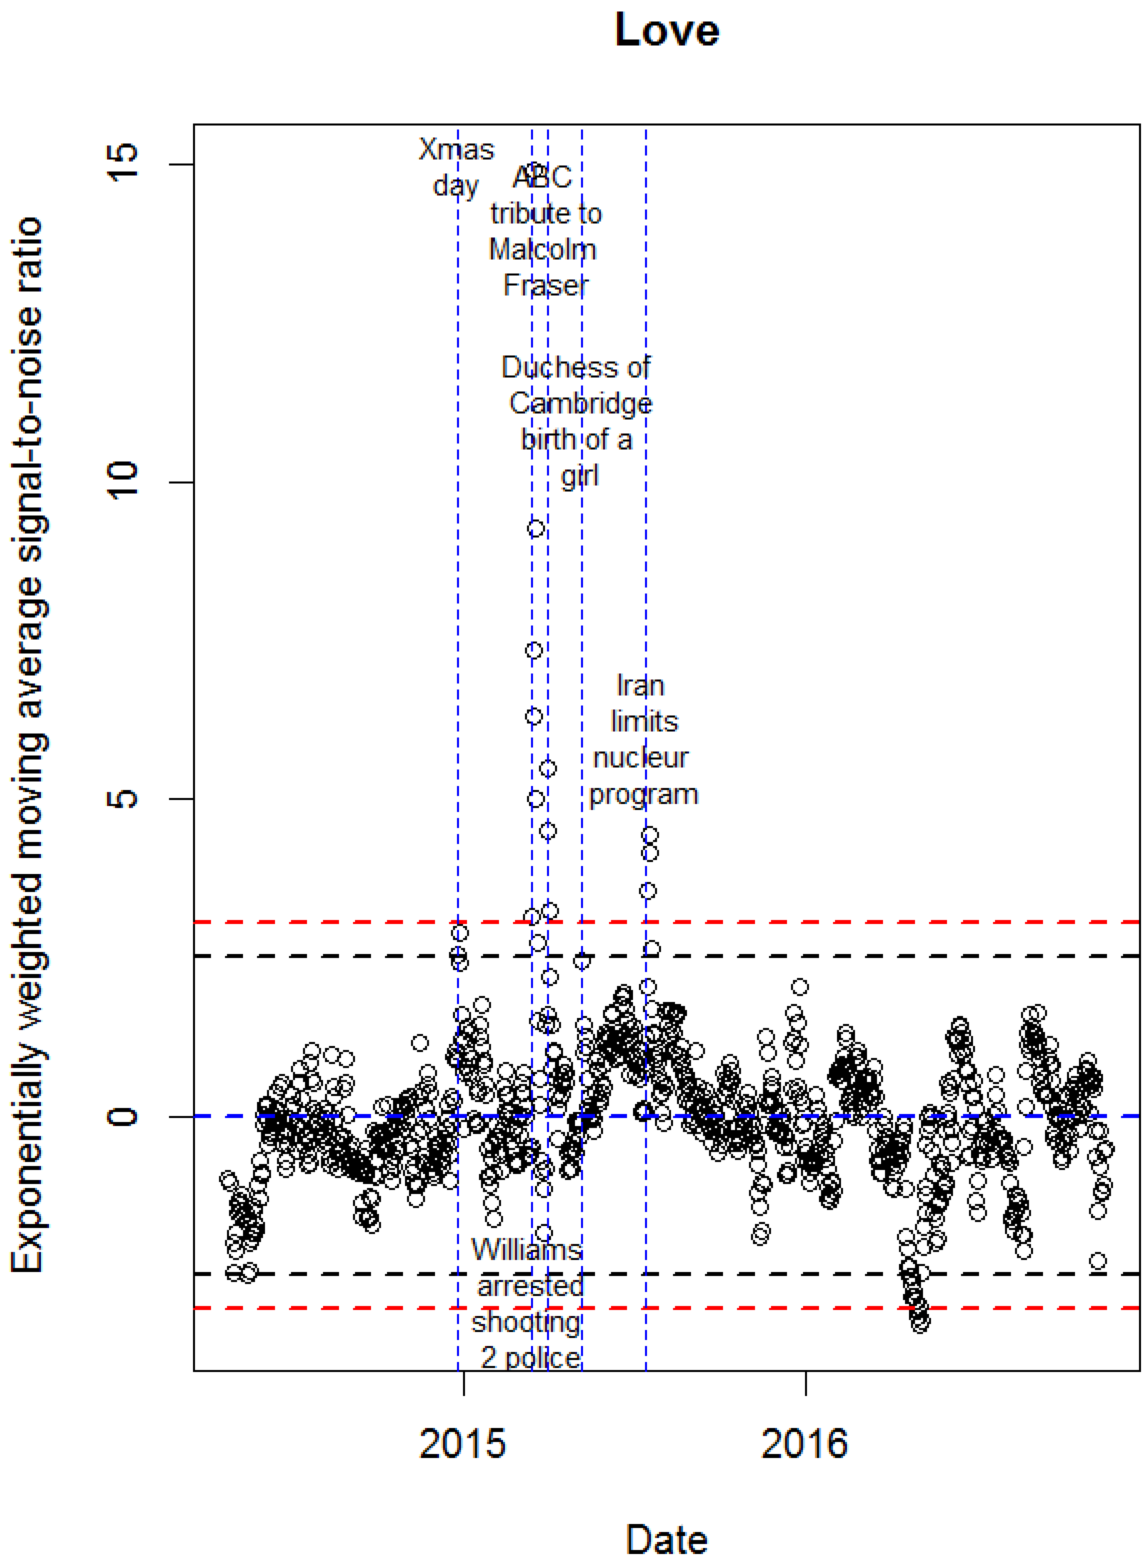
Love
Xmas day
ABC tribute to Malcolm Fraser
Duchess of Cambridge birth of a girl
Iran limits nucleur program
Williams arrested shooting 2 police
Exponentially weighted moving average signal-to-noise ratio
15
10
5
0
2015
2016
Date

Figure 9:

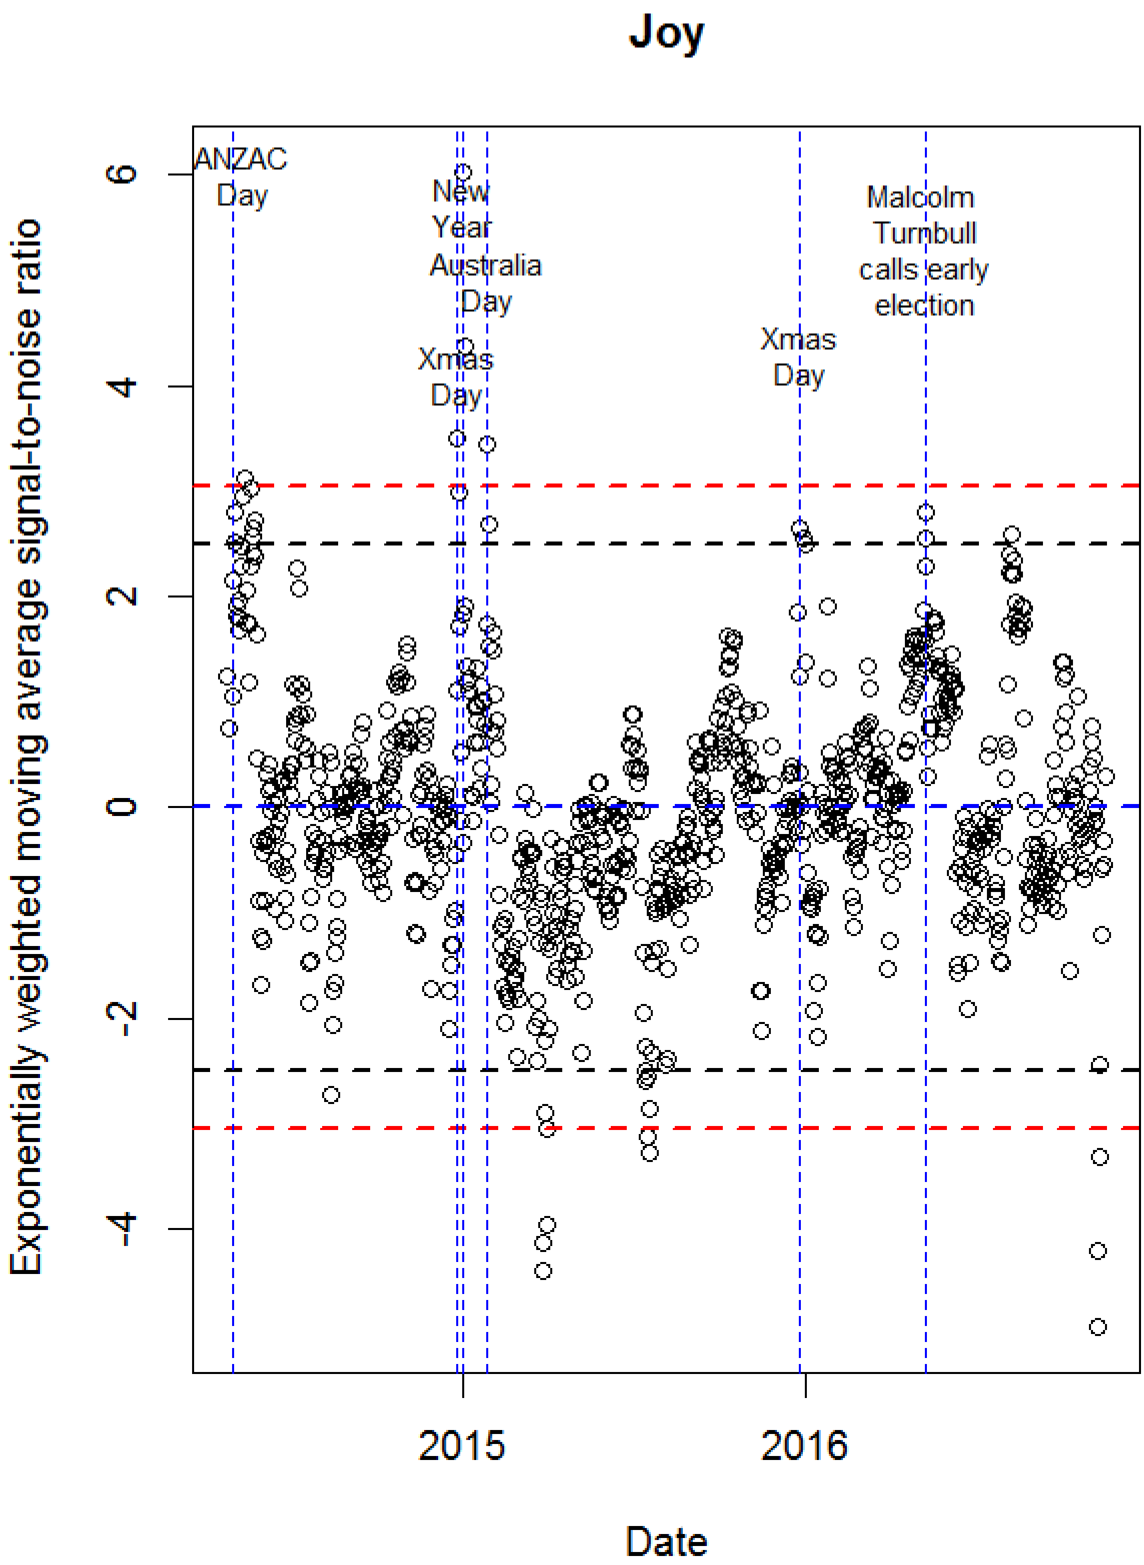

12

Figure 10: Fig. 12 :

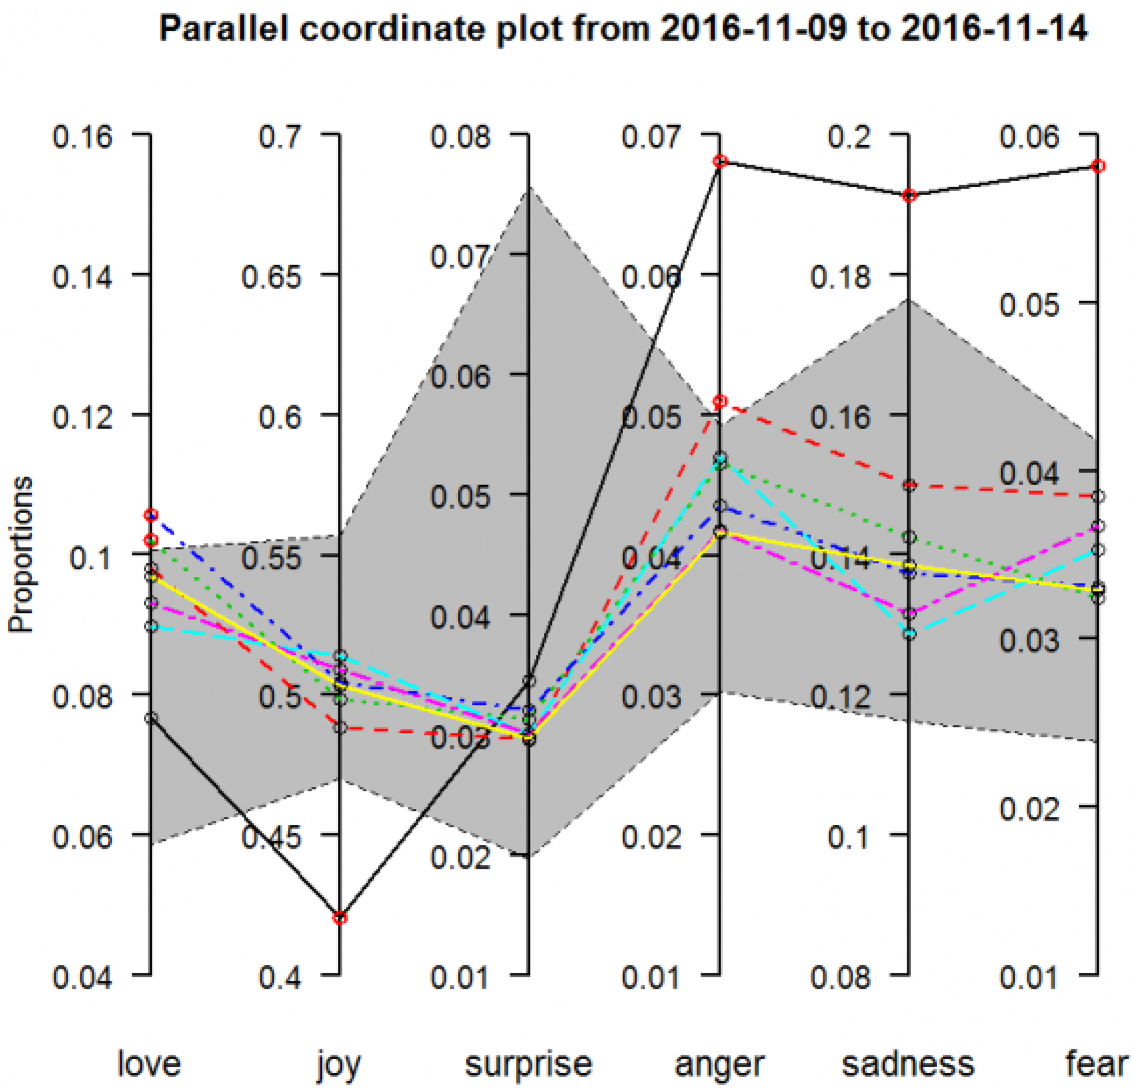

13

Figure 11: Fig. 13 :

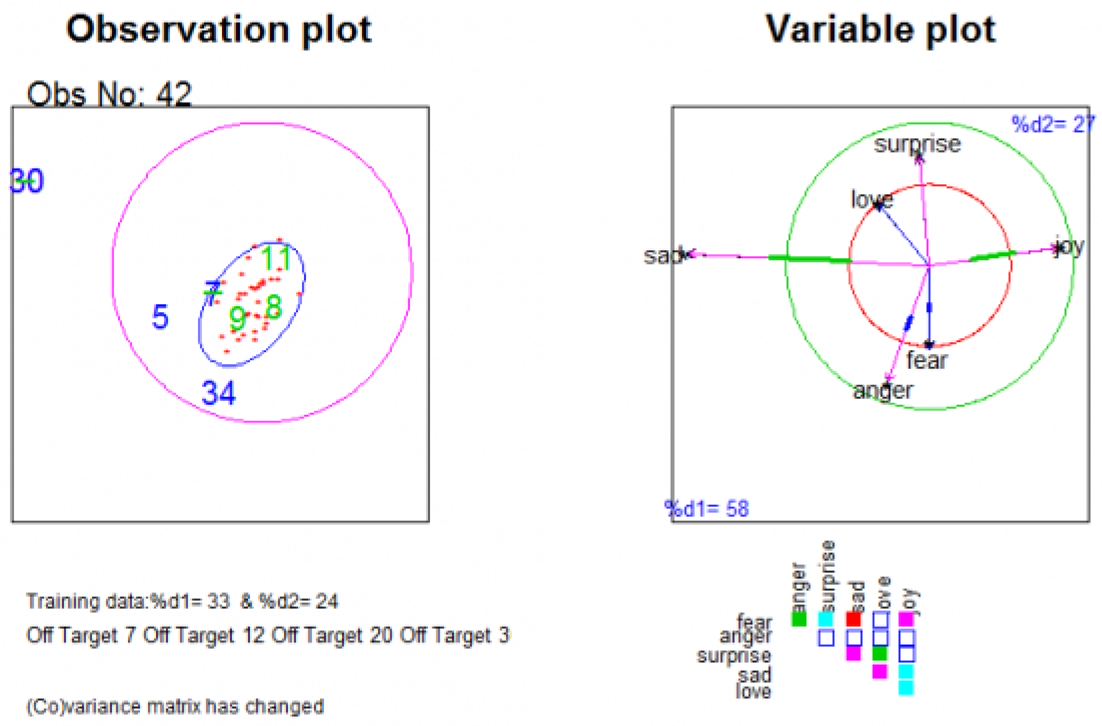

14

Figure 12: Fig. 14 :

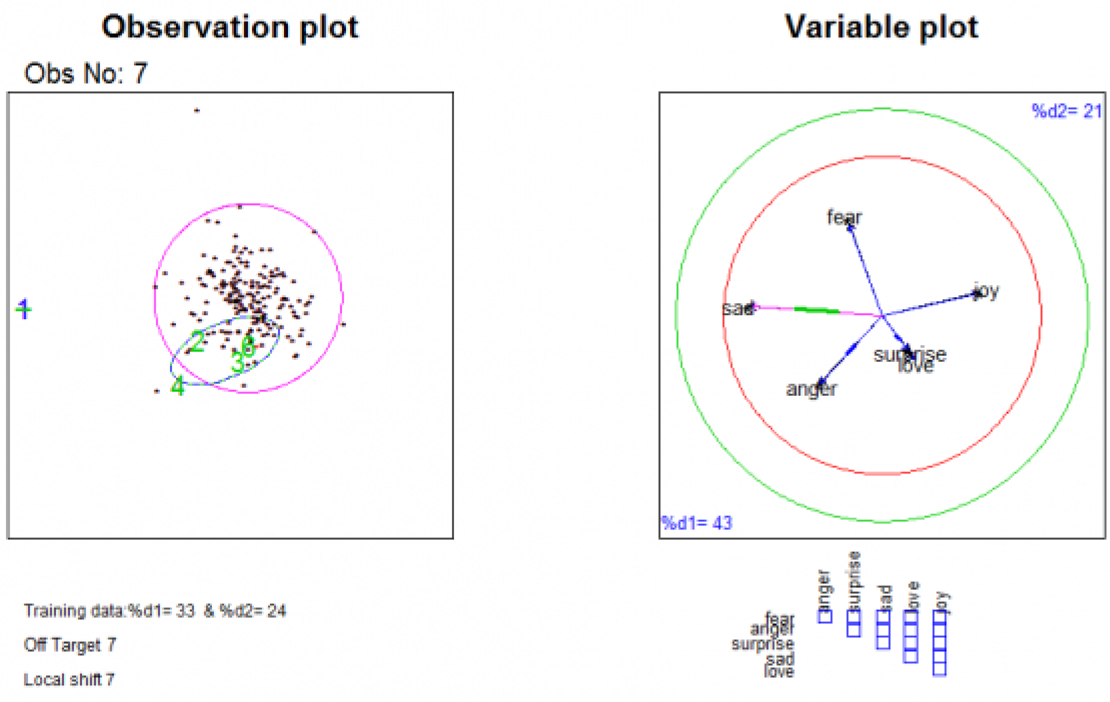

16

Figure 13: Fig. 16 :

Figure 14:

[Paris et al. ()] 'Exploring emotions in social media'. Cécile Paris , Christensen , Helen , Philip Batterham , O' Dea , Bridianne . *the Proceedings of the IEEE International Conference on Collaboration and Internet Computing (CIC). Oct 27-31 st , Hang Zhou*, (China) 2015.

[Sparks et al. ()] 'Multivariate process monitoring using the dynamic biplot'. R Sparks , A Adolphson , A Phatak . *International Statistical Review* 1997. 65 (3) p. .

[Larsen et al. ()] 'We Feel: Mapping emotions on Twitter'. M Larsen , T Boonstra , P Batterham , B O'dea , C Paris , H Christensen . 10.1109/JBHI.2015.2403839. *IEEE Journal of Biomedical and Health Informatics (JBHI)* 2168-2194. 2015.