Lansdall-Welfare, T., Dzogang, F., & Cristianini, N. (2017). Change-Point Analysis of the Public Mood in UK Twitter during the Brexit Referendum. In *2016 IEEE 16th International Conference on Data Mining Workshops (ICDMW 2016)* Institute of Electrical and Electronics Engineers (IEEE).
<https://doi.org/10.1109/ICDMW.2016.0068>

Peer reviewed version

Link to published version (if available):
[10.1109/ICDMW.2016.0068](https://doi.org/10.1109/ICDMW.2016.0068)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the author accepted manuscript (AAM). The final published version (version of record) is available online via IEEE at <http://ieeexplore.ieee.org/document/7836699/>. Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

Change-point Analysis of the Public Mood in UK Twitter during the Brexit Referendum

Thomas Lansdall-Welfare, Fabon Dzogang and Nello Cristianini

Intelligent Systems Laboratory

University of Bristol

Bristol

United Kingdom

Email: firstname.lastname@bristol.ac.uk

Abstract—We study the changes in public mood within the contents of Twitter in the UK, in the days before and after the Brexit referendum. We measure the levels of anxiety, anger, sadness, negative affect and positive affect in various geographic regions of the UK, at hourly intervals. We analyse these affect time series’ by looking for change-points common to all five components, locating points of simultaneous change in the multivariate series using the fast group LARS algorithm, originally developed for bioinformatics applications. We find that there are three key times in the period leading up to and including the EU referendum. In each case, we find that the public mood is characterised by an increase in negative affect, anger, anxiety and sadness, with a corresponding drop in positive affect. The hour by hour evolution of public mood in the hours leading up to and following the closure of the polls is further analysed in conjunction with the GBP/EUR exchange rate, finding four change-points in the hours following the vote, and significant correlation between the exchange rate and the affect components tested.

Index Terms—Public Mood, Social Media, Politics, Brexit, Change-point Analysis, Information Fusion, Big Data, Multivariate Time series.

I. INTRODUCTION

Public mood has been shown to play a role in how individuals process information and form political opinions, with affective experiences influencing political reasoning [1]. However, assessing public mood on a large scale using traditional methods can be prone to error from a range of sources including the wording of questions [2], individuals using their momentary affective state to judge their overall state [3] and even the characteristics and background of the interviewer [4]. To overcome many of these issues, use of social media as an alternative method for assessing the public mood has been demonstrated in previous studies (e.g. [5]–[9]).

The causes behind changes in public mood can be difficult to explicitly capture, due in part to its diffuse nature [10]. Previous work has gone some way in explaining the changes in public mood as measured using social media, finding circadian and seasonal patterns of affect [5], [6], along with public mood changing in response to specific real-world events [7], [8].

In this study, we are interested in investigating changes in public mood through analysis of simultaneous and sudden changes in five affect components, and the events that triggered them. We make use of the fast group LARS algorithm [11],

that detects shared change-points across several time series’ at once. These specific points in time, where many time series’ change together, have the potential of signalling specific real-world events that explain the variation in the public mood.

We find that there are three key times in the period leading up to and including the European Union (EU) referendum, coinciding with the football violence in Marseille between English and Russian fans and the Orlando nightclub shooting, the murder of Labour MP Jo Cox, and the results of the EU referendum itself. In each of these cases, the public mood is characterised by a decrease in positive affect, and an increase in negative affect, anger, anxiety and sadness, with the reaction corresponding to when the outcome of the referendum result became clear causing the largest negative change in public mood.

Furthermore, we find that analysing the affect components in different geographical regions of the United Kingdom shows a robust signal, with each region following a very similar trajectory over the period, and that the hour by hour evolution of public mood in the 48 hours starting on the day of the referendum significantly correlates with the GBP/EUR exchange rate for the five affect components used in this study.

II. METHODS

A. Data collection

We gathered social media data from Twitter, an online platform that allows users to publish brief textual communications (tweets) of up to 140 characters, which are publicly visible and available via their application programming interface (API). Using the Twitter API, we collected over 10 million tweets during a period of 30 days between 1st June 2016 and 30th June 2016, querying for tweets geo-located to within 10km of any of the 54 largest urban centres in the United Kingdom, without specifying any keywords or hashtags. For each tweet, we collected the anonymised textual content, a collection date and time, and information about the location from where the tweet was collected (one of the 54 urban centres).

Tweets were preprocessed into their constituent tokens using a tokenizer designed specifically for Twitter text [12]. Tokens representing hyperlinks, mentions and hashtags were discarded, along with tokens containing only special characters (e.g. emoticons).

Additional data on the exchange rate between Pound Sterling and the Euro (GBP/EUR) was collected from the web. In this study we use the opening price of the Pound against the Euro in hourly intervals, converted from the Euro against the Pound¹.

B. Token time series generation

Each token found within the tweets was converted to an hourly time series representation, representing how often it was present within the collected tweets over the 30 days investigated in the study. This consisted of two steps, first counting the raw number of times each token occurs per hour, followed by computing the relative frequency of each token, allowing us to perform a fair comparison of the usage of tokens across hours with differing numbers of tweets.

Using the map-reduce framework, commonly used in big data applications, we counted the number of times each token occurred within tweets collected within each hour, giving us a 720-length (24 hours $\times$ 30 days) time series for the raw count of each token. The hourly volume of tweet tokens was then computed by summing over each of the individual token time series, giving us a single hourly time series representing the total number of tokens across all tweet published in each hour.

Before normalising the raw token count time series' with the hourly volume time series to obtain each token's relative frequency time series, we applied a three-hour centred moving average to both the raw counts and the hourly volume to improve the estimation of each token's frequency. The moving average is applied to ensure that we have enough statistics to estimate the relative frequency of each token, where many tokens can be rare, due to Zipf's law [13], or due to low volume hours.

C. Measuring public mood

As has become standard in many sentiment analysis studies [5], [9], [14], we take a lexicon-based approach to sentiment and affect analysis in text. We measure five components of public mood using the Positive Affect (PA), Negative Affect (NA), Anger, Anxiousness and Sadness lexica contained within the Linguistic Inquiry and Word Count (LIWC) [15].

The LIWC lexicon contains many word lists that measure different dimensions of psychological and behavioural characteristics in text, including the five previously mentioned components we consider in this study. The lexicon was designed to be applied to a wide range of different texts, including transcribed every day speech and email, making it suitable for application in social media domains. Furthermore, the lists were validated by independent judges, and found to have high levels (0.88 and 0.97 respectively) of sensitivity and specificity for all emotional expression words [16].

For each of the five affective components we wished to measure within Twitter, we extracted the list of tokens from LIWC and retrieved the set of corresponding standardized relative frequency time series' for all related tokens. The set

of time series' were then averaged across all tokens within an affect component, resulting in a single overall series for each of the five affect components: PA, NA, Anger, Anxiety and Sadness.

Due to the highly circadian nature of affect [5], [6], we apply a smoothing function to the affect time series' to account for the natural daily fluctuations, allowing us to study changes in overall affect which are not explained by the typical circadian pattern. Each time series is finally standardized by subtracting the mean and normalising by the standard deviation to obtain comparable time series' with zero mean and unit variance. For analysing a shorter time series, we alternatively detrend the time series instead of smoothing as discussed in Section II-E.

D. Decomposition into regional time series'

Further to the overall or 'national' level of the five affect components calculated as above, each affect component was also calculated separately for the following twelve regions of the United Kingdom: North East, North West, Yorkshire, East Midlands, West Midlands, East of England, South East, South West, London, Northern Ireland, Scotland and Wales.

This was performed by generating region-specific token time series' for each of the twelve regions, using all tweets published from locations falling within that region of the UK. The same procedure for measuring public mood was then followed for each region, resulting in 60 affect time series' (12 regions $\times$ 5 affect components).

E. Analysing hourly changes

It can also be of interest to "zoom in" to a shorter segment of a time series and more closely analyse the data found within a particular period at a finer resolution. However, due to the smoothing applied to account for the circadian pattern, it becomes difficult to clearly identify the exact time of changes in periods shorter than the smoothing window. We therefore use an alternative method for removing the circadian pattern for short time periods.

For a given segment of a time series, we remove the circadian pattern by performing a detrending step used in signal processing applications [17]. Specifically, we calculate the median circadian pattern over the three week-days previous to the segment of interest on the affect time series without any smoothing, before removing this trend from all days within the segment, resulting in a detrended segment with the circadian pattern removed where no smoothing has been applied.

F. Group change-point detection

We analyse the affect time series' by looking for change-points common to the five components, where our aim is to locate points of simultaneous change in the multivariate series, rather than for each component separately. Using the fast group LARS algorithm [11], a tool originally designed for the analysis of genomic profiles in bioinformatics, we compute a piecewise constant approximation of each affect component endowed with the property of common connected regions.

¹Historical Forex data available from <http://www.histdata.com/download-free-forex-historical-data/?/metatrader/1-minute-bar-quotes/eurgbp/2016/6>

Fig. 1. Standardized scores for the national level of five affect components during June 2016 in the United Kingdom. Identified change-points are indicated with vertical dashed lines, with the piecewise constant value between change-points indicated with a solid black line between change-points. Change-points within the first and last 24 hours were discounted due to the effect of smoothing with a centred moving average at the boundary.

Specifically, the multivariate affect series' Y are reconstructed under the constraint of sparse successive differences cancelled groupwise in time [11], [18]. This can be formally expressed as the convex optimization problem:

$$\min_{U \in \mathbb{R}^{n \times p}} \frac{1}{2} \|Y - U\|_2^2 + \lambda \sum_{i=1}^{n-1} \epsilon_i \|U_{i+1, \bullet} - U_{i, \bullet}\|_2, \quad (1)$$

where $U_{i, \bullet}$ denotes the i -th row of U , p is the number of affect time series', n is the length of the affect time series', λ penalises the group total variation and $\epsilon_i > 0$ is a position-dependent correction used to alleviate some boundary effects. The solution to (1) is then further fine-tuned using dynamic programming as described in [11].

III. RESULTS

We focus our analysis on the time around the United Kingdom's referendum on remaining a member of, or leaving the EU, commonly referred to as Brexit, which took place on the 23rd June 2016, with a final outcome of 51.9% voting to leave the EU.

A. Public mood at the national level

Figure 1 shows the five affect components computed at the national level for the 30 days in June 2016. Change-points identified using the fast group LARS algorithm are indicated by dashed vertical lines, while the piecewise constant between change-points is indicated in black, showing the mean level of each affect component between the change-points. In total, 13 change-points during the 30 days are found, corresponding with the times when all five affect components changed the most simultaneously. Change-points within the first and last 24 hours were discounted due to the effect of smoothing with a centred moving average at the boundary.

Here we analyse the main change-points which correspond with the three most striking peaks occurring jointly in the negative components (NA, anger, anxiety and sadness), along with nadirs in the PA time series. These change-points highlight the periods between 11th and 13th June, the 17th June and between the 23rd and 25th June 2016 as being those times of greatest change across all affect components.

1) *11th - 13th June 2016*: On the 11th June 2016, during the 2016 UEFA European Championship taking place in

Fig. 2. Standardized scores for the regional level of five affect components during June 2016 for the 12 regions of the United Kingdom. Identified change-points are indicated with vertical dashed lines, with the piecewise constant value between change-points indicated with a solid black line between change-points. Change-points within the first and last 24 hours were discounted due to the effect of smoothing with a centred moving average at the boundary.

France, violence broke out between football supporters at the close of the England vs. Russia game held at the Stade Vélodrome in Marseille. The clashes between fans were not limited to within the stadium, with at least 20 supporters injured before the game, and on-going aggression following the game [19], along with a strong media coverage of the events at the time.

In Fig. 1 we can see change-points corresponding with the early afternoon on the day of the game, with a further change-point identified in the early hours of 12th June, before a final change-point on the afternoon of the 13th June when the levels of NA and anger began to subside. We can clearly see that the reaction on social media at the time was characterised by increases in NA, anger and anxiety, with a marked decrease in PA, while sadness shows a much smaller increase during this period.

Events unfolding in the early hours of the 12th June in the United States, where a nightclub in Orlando was attacked by a lone gunman resulting in 49 people being killed, offer further explanation to the sudden changes found within the affect components in this period.

2) *17th June 2016:* We identified a change-point on the 17th June 2016 when the negative components decrease, and PA increases, following a peak the day before in all negative components. This peak of anger, anxiety, sadness and NA corresponds with the murder of Labour MP Jo Cox after she was shot and stabbed after holding a constituency meeting on the afternoon of the 16th June 2016 [20]. It is remarkable to note that this tragic event came at a time of heightened negative feeling in social media following the football violence in France, and also how quickly the overall public mood returned to previous levels following. It should be clearly

Fig. 3. Standardized scores for the detrended national level of five affect components and the exchange rate between GBP and EUR during the 48 hours starting from the 23rd June 2016, covering both the period before and after the referendum results being announced. Identified change-points are indicated with vertical dashed lines.

stated however that this does not imply that such an event took place *because* of the increased levels of negative feelings that were taking place in the country at the time.

3) *23rd - 25th June 2016*: On the 23rd June 2016, the United Kingdom voted on whether to remain a member of the EU, or to leave, following a somewhat controversial campaign on both sides that were criticised for being ‘highly misleading to the electorate’ [21]. The outcome of the referendum came as a surprise to many, with the final YouGov poll before the result incorrectly giving Remain a four point lead [22], bookmakers predicting an 86.29% chance of a Remain victory just hours before the polls closed [23], and Nigel Farage, a leading proponent of the UK leaving the EU, conceding as the polling stations closed that it “looks like Remain will edge it” [24].

However, as the results started to be announced, it quickly became clear that the result was pointing towards a victory for Leave. Figure 1 shows a sharp increase in the NA, anger, anxiety and sadness, and a drop in PA following the closure of the polls and the results beginning to be announced. We investigate this further in Section III-C where we analyse the hourly changes in the affect components at a finer resolution.

B. Public mood at the regional level

Figure 2 shows the five affect components computed at the regional level for the 30 days in June 2016 in the 12 regions of the United Kingdom. Change-points are identified using the fast group LARS algorithm on the 60 time series’, and are indicated by dashed vertical lines, while the piecewise constant between change-points is indicated in black, showing the mean level over all regions for each affect component between the change-points. In total, 13 change-points during the 30 days are found, with only slight differences with those found when computed at the national level. Change-points within the first and last 24 hours were again discounted due to the effect of smoothing with a centred moving average at the boundary.

The differences in change-points at the regional level include: an additional minor change-point on the 2nd June 2016 with PA increasing slightly, while some negative components decreased; splitting the minor change-point on the 1st June

2016 into two which are a couple of hours apart; and collapsing the major change-point following the referendum result on the 25th June 2016 into a single change-point. We also see that the period of change from the 11th to the 13th June is extended to the 15th June when computed at the regional level, highlighting that different regions expressed their feelings for slightly different lengths of time.

In some instances, we can see how individual regions deviated from the rest of the country during the main change-points. For example, in Fig. 2 we observe an exaggerated NA and anger response towards the end of 16th June 2016 in Wales. While this is during the peak of all negative components following the murder of Jo Cox, the specific region response is perhaps better explained by the European Championship match on the same day between England and Wales which lead to Wales’ defeat.

On the whole however, we found that the major change-points identified at both the national and regional level are stable, and that the affective response from each region was surprisingly uniform, following a very similar trajectory over the 30 days.

C. Closer inspection of the referendum reaction

Finally, we wished to more closely inspect the changes happening in the lead up to and following the referendum, zooming into the 48 hours starting from midnight on the 23rd June 2016 and finishing on midnight of the 25th June 2016, covering the hours when the polls were open, and the reaction to the results being announced in the following 24 hours after polling closed.

Figure 3 shows the national level of the five affect components over 48 hours starting from midnight on the 23rd June 2016, re-standardized within the 48-hour window, and displayed with the standardized exchange rate between the Pound and Euro during the period following the procedure in Section II-E.

Calculating change-points for the 48-hour period, we found that change-points occur around 1am and 5am in the early morning of the 24th June as the results for the different districts are being announced, later in the morning at 11am, then again at 5pm in the afternoon. The final change-point

TABLE I
PEARSON CORRELATION BETWEEN THE GBP/EUR EXCHANGE RATE AND
THE FIVE AFFECT COMPONENTS IN THE 48 HOUR PERIOD.

	PA	NA	Anger	Anxiety	Sadness
GBP/EUR	0.48	-0.84	-0.73	-0.42	-0.89

identified at 10pm corresponds with the end of the tested period for change-points, after which the forex data was unavailable due to the markets closing.

We additionally calculated the correlation between the exchange rate and the five affect components as shown in Table I, finding that while PA shows a positive correlation with the exchange rate, a stronger anti-correlation is found for NA, anger and sadness, with sadness measured in Twitter explaining the greatest variance in the exchange rate out of the five affect components during these 48 hours. All correlation coefficients were found to be statistically significant at the 5% level, assessed using the Student's t-test and corrected for multiple testing using the Bonferroni correction.

IV. DISCUSSION

Fluctuations in collective public mood are due to a multitude of competing effects, some of which are seasonal and predictable [5], while others are driven by external events [7].

In this study, we have demonstrated a first case in which we attempt to explain the variability of public mood through simultaneous multiple change-point analysis. When combined with other sources of variance, this approach can reduce the number of unexplained movements in public mood. This has implications for the political sciences, where understanding changes in public mood can help elucidate the link between affective experiences and the formation of political opinions. The constant monitoring of public mood in both conventional and social media has the potential of providing real insight into how events and policies influence public attitudes.

More generally, the methodology outlined in this study is general, and can be transferred to many other domains where simultaneous change-points in multivariate series' need to be detected. This provides a succinct way to perform information fusion across data coming from disparate sources, as evidenced by the group change-points found in both the public mood and exchange rate, and the correlation found between them in this study.

ACKNOWLEDGMENTS

The authors would like to thank Jean-Philippe Vert and Kevin Bleakley for making their group fused Lasso code available. This study was funded the ERC Advanced Grant ThinkBIG.

REFERENCES

- [1] W. M. Rahn, "Affect as information: The role of public mood in political reasoning," *Elements of reason: Cognition, choice, and the bounds of rationality*, pp. 130–50, 2000.
- [2] S. L. B. Payne, *The Art of Asking Questions: Studies in Public Opinion*, 3. Princeton University Press, 2014, vol. 3.
- [3] N. Schwarz and G. L. Clore, "Mood, misattribution, and judgments of well-being: Informative and directive functions of affective states," *Journal of personality and social psychology*, vol. 45, no. 3, p. 513, 1983.
- [4] I. C. Choi and G. W. Comstock, "Interviewer effect on responses to a questionnaire relating to mood," *American journal of epidemiology*, vol. 101, no. 1, pp. 84–92, 1975.
- [5] S. A. Golder and M. W. Macy, "Diurnal and seasonal mood vary with work, sleep, and daylength across diverse cultures," *Science*, vol. 333, no. 6051, pp. 1878–1881, 2011.
- [6] V. Lampos, T. Lansdall-Welfare, R. Araya, and N. Cristianini, "Analysing mood patterns in the united kingdom through twitter content," *arXiv preprint arXiv:1304.5507*, 2013.
- [7] T. Lansdall-Welfare, V. Lampos, and N. Cristianini, "Effects of the recession on public mood in the uk," in *Proceedings of the 21st International Conference on World Wide Web*. ACM, 2012, pp. 1221–1226.
- [8] V. D. Nguyen, B. Varghese, and A. Barker, "The royal birth of 2013: Analysing and visualising public sentiment in the uk using twitter," in *Big Data, 2013 IEEE International Conference on*. IEEE, 2013, pp. 46–54.
- [9] J. Bollen, H. Mao, and A. Pepe, "Modeling public mood and emotion: Twitter sentiment and socio-economic phenomena," *ICWSM*, vol. 11, pp. 450–453, 2011.
- [10] W. M. Rahn, B. Kroeger, and C. M. Kite, "A framework for the study of public mood," *Political Psychology*, pp. 29–58, 1996.
- [11] K. Bleakley and J.-P. Vert, "The group fused lasso for multiple change-point detection," *arXiv preprint arXiv:1106.4199*, 2011.
- [12] K. Gimpel, N. Schneider, B. O'Connor, D. Das, D. Mills, J. Eisenstein, M. Heilman, D. Yogatama, J. Flanagan, and N. A. Smith, "Part-of-speech tagging for twitter: Annotation, features, and experiments," in *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies: short papers-Volume 2*. Association for Computational Linguistics, 2011, pp. 42–47.
- [13] G. Zipf, "Selective Studies and the Principle of Relative Frequency in Language, volume 23," 1932.
- [14] A. Tumasjan, T. O. Sprenger, P. G. Sandner, and I. M. Welp, "Predicting elections with twitter: What 140 characters reveal about political sentiment," *ICWSM*, vol. 10, pp. 178–185, 2010.
- [15] J. W. Pennebaker, M. E. Francis, and R. J. Booth, "Linguistic inquiry and word count: Liwc 2001," *Mahway: Lawrence Erlbaum Associates*, vol. 71, 2001.
- [16] E. O. Bantum and J. E. Owen, "Evaluating the validity of computerized content analysis programs for identification of emotional expression in cancer narratives," *Psychological assessment*, vol. 21, no. 1, p. 79, 2009.
- [17] P. Flandrin, P. Goncalves, and G. Rilling, "Detrending and denoising with empirical mode decompositions," in *Signal Processing Conference, 2004 12th European*. IEEE, 2004, pp. 1581–1584.
- [18] R. Tibshirani, M. Saunders, S. Rosset, J. Zhu, and K. Knight, "Sparsity and smoothness via the fused lasso," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 67, no. 1, pp. 91–108, 2005.
- [19] T. Ough, T. Morgan, and C. Criddle. (2016, June) Euro 2016: Uefa threatens to disqualify england and russia amid german and ukraine violence in lille. The Telegraph. [Online]. Available: <http://www.telegraph.co.uk/football/2016/06/12/euro-2016-uefa-investigate-violence-between-england-and-russia-f/>
- [20] R. Booth, V. Dodd, and N. Parveen. (2016, June) Labour mp jo cox dies after being shot and stabbed. The Guardian. [Online]. Available: <https://www.theguardian.com/uk-news/2016/jun/16/labour-mp-jo-cox-shot-in-west-yorkshire>
- [21] H. of Commons Treasury Committee, "The economic and financial costs and benefits of the UKs EU membership," 2016. [Online]. Available: <http://www.publications.parliament.uk/pa/cm201617/cmselect/cmtreasy/122/122.pdf>
- [22] J. Twyman. (2016, June) Yougov on the day poll: Remain 52[Online]. Available: <https://yougov.co.uk/news/2016/06/23/yougov-day-poll/>
- [23] M. Dent. (2016, June) Referendum odds — eu brexit. PoliticalOdds.bet. [Online]. Available: <http://politicalodds.bet/eu-referendum>
- [24] C. Cooper and K. Forster. (2016, June) Eu referendum: Nigel Farage says it 'looks like remain will edge it' as polls close. The Independent. [Online]. Available: <http://www.independent.co.uk/news/uk/home-news/eu-referendum-nigel-farage-remain-edge-it-brexit-ukip-a7098526.html>