# The Effect of the Recession on Collective Mood in the UK

**Thomas Lansdall-Welfare, Vasileios Lampos, Nello Cristianini**
**Intelligent Systems Laboratory, University of Bristol**

## Abstract

We report on the analysis of 484 million tweets collected in the United Kingdom between July 2009 and January 2012, a period marked by economic downturn and some social tensions. Our analysis of sentiment expressed in that time series reveals various change points in collective mood, most noticeably two negative shifts corresponding to the announcement of spending cuts and to the riots of August 2011. We propose that constant nowcasting of certain collective properties of society is possible by monitoring the contents of social media.

## Introduction

Social media allows for the easy gathering of large amounts of data generated by the public while communicating with each other. This can give social scientists – but also policy makers and decision makers - a tool to access public opinion and public sentiment, a task so far accomplished only by opinion polls. Large-scale text-mining and data-mining infrastructures need to be developed in order to exploit this resource. Twitter in particular gives researchers access to real-time data that is suitable for the analysis of public sentiment on a Big-Data scale, but the same techniques could be applied to any other user-generated textual-stream, if it was available.

This type of analysis is of interest because it avoids self-reporting and opinion-polling, and therefore opens the possibility to access much larger populations. It also allows for instantaneous gauging of social trends, a task that we call "nowcasting", in contrast to "forecasting". Measuring the state of society at any given time is a complex operation that usually cannot be done in real-time.

For all its appeal, the real-time detection of social trends via the analysis of social media content also presents various possible drawbacks. For a start, this can only be accomplished with text mining technologies, which are less accurate than human assessment, but can be applied to vast amounts of data. Also the population that is assessed is necessarily that of Twitter users, which is a biased subsample of the general population. So, particular care needs to be taken when extracting information, but also when reporting it.

Recent studies in this growing field range from trying to predict the stock market (Bollen, Mao and Zeng, 2010) and determining seasonal changes in mood (Golder and Macy, 2011) to gauging political opinion (O'Connor, Balasubramanyan, Routledge and Smith, 2010) or current levels of flu in the population (Lampos, De Bie and Cristianini, 2010). This shows a diverse array of potential problems which can be approached in a real-time data driven way.

In this study we focus on measuring the mood (and changes thereof), using standard tools for mood detection (Golder and Macy, 2011; Dodds, Harris, Kloumann, Bliss and Danforth, 2011), of a large sample of the UK population. In particular, we analyse a collection of 484 million tweets generated by more than 9.8 million users from the United Kingdom between July 2009 and January 2012, a period marked by economic downturn and some social tensions.

Our findings (Lansdall-Welfare, Lampos, Cristianini, 2012) present intriguing patterns that can be explained in terms of events and social changes. Among them, we see that a significant increase in negative mood indicators coincide with the announcement of the cuts to public spending by the government, and that this effect is still lasting. We also detect events such as the riots of summer 2011, as well as a possible calming effect coinciding with the run up to the royal wedding.

**Methodology**

Social media is ubiquitously accessible, enabling individuals and communities to take part in interactive communication and dialogue on a large scale. In particular, we focus on Twitter which provides its users with a platform for publishing publicly available textual communications of up to 140 characters. Perhaps due to their brief nature, the communications – or 'tweets' as they are known – tend to be in-the-moment expressions of the user's current experiences and feelings, making them an ideal candidate for easily capturing the state of an individual's mood.

In this study, we collected approximately 484 million tweets between 1st July 2009 and 19th January 2012 authored by 9,812,618 users in the UK. Collection took place by retrieving the 100 most recent tweets every 3-5 minutes for each of the 54 most populous urban centres in the UK, geo-located to within a 10km range of the centres. Each tweet, once retrieved, was stemmed to remove word inflection and plurality by applying Porter's Algorithm (Porter, 1993) before using a text analysis tool to assess the emotional content.

There are standard methods in text analysis for detecting sentiment known in the trade as citation-sentiment analysis: they are commonly used in marketing research to determine the opinion of users of a certain product or viewers of a television show. For our purposes, each of the four basic emotions (fear, joy, anger, sadness) is associated with a list of words which have been generated by a combination of manual and automated methods, and successively benchmarked on a test set. These word lists were based on extracts taken *a priori* from the tool 'WordNet Affect' (Strapparava and Valitutti, 2004). The lists were also stemmed to remove word inflection and filtered in a way to only keep single words. After this pre-processing, the word lists contained 146 anger words, 92 fear words, 224 joy words and 115 sadness words.

A daily representation of each emotion was calculated by first scoring each word in the corresponding word list as the proportion of that day's tweets containing it. This quantity is then averaged over all words in the word list to give an overall valence for that particular emotion on that day. This method is based on the assumption that the average frequency of a word indicates its importance; words with higher frequencies will have a larger impact on the valence of each emotion. Repeating this process for each day in the dataset gives a timeline of the changes in public mood over the 31 month period.

**Analysis**

The first part of our analysis focuses on verifying that word-counting methods can provide a reasonable approach to detecting a signal of collective mood, captured from public Twitter messages.

We found that all four of the key emotions change over time, in a manner that is partly predictable (or at least interpretable), based upon the time of year and current events at the time. For example, in Figure 1 we can see that each year there was a peak of joy around Christmas – surely due to season's greetings - and another periodic peak, this time of fear, around Halloween as seen in Figure 2. Again, one would imagine this is probably due to increased usage of certain words such as 'scary' found in the word lists. Of course, we do not expect that a high frequency of the word 'happy' necessarily signifies happier mood in the population. Our measures of mood are not perfect, but these effects could be filtered away by more sophisticated tools designed to ignore conventional expressions such as 'Merry Christmas' or 'Happy New Year'. However, it is a remarkable observation that reliably similar values can be found for certain days across the different years, suggesting that we have reduced statistical errors to a low level.

The main part of the analysis looks for change-points occurring in the public's mood, times when large shifts in the collective sentiment of the population can be detected. Furthermore, we can show that the change-points are real and statistically significant, and that in some cases the effects are still lasting.

Change-points are detected by looking for abrupt changes in the mean value of the mood signal. We found abrupt changes in the mean by using a 100-day moving window (50 days either side), calculating the difference in mean before and after each day. This measure is used to show the rate of mood change for a given day, and can informally be thought of as an indication of the derivative of the overall mood level.

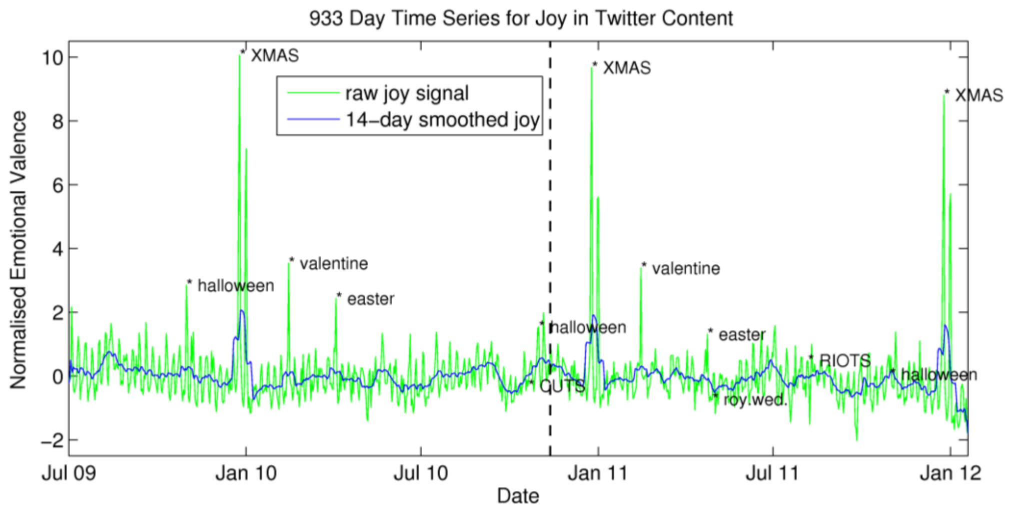


**Figure 1: Plot showing a time series for joy words found in the contents of Twitter in the United Kingdom between 1st July 2009 and 19th January 2012.**

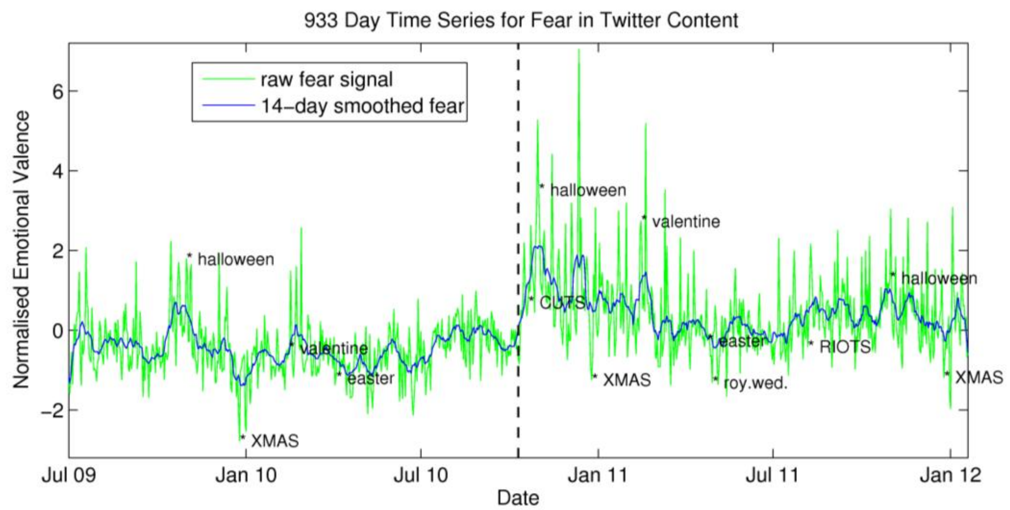


**Figure 2: Plot showing a time series for fear words found in the contents of Twitter in the United Kingdom between 1st July 2009 and 19th January 2012.**

Two such change points were detected using this measure, the first of which was in October 2010, when Prime Minister Gordon Brown announced massive cuts in public spending. Although the announcement did not come as a surprise to the public, with extensive news media coverage of the upcoming event, it was a clear change-point, further validated by a statistical test, and it was - if you like - the moment that people realised that austerity was not just for others; it would be affecting their own lives too. The effects of that major

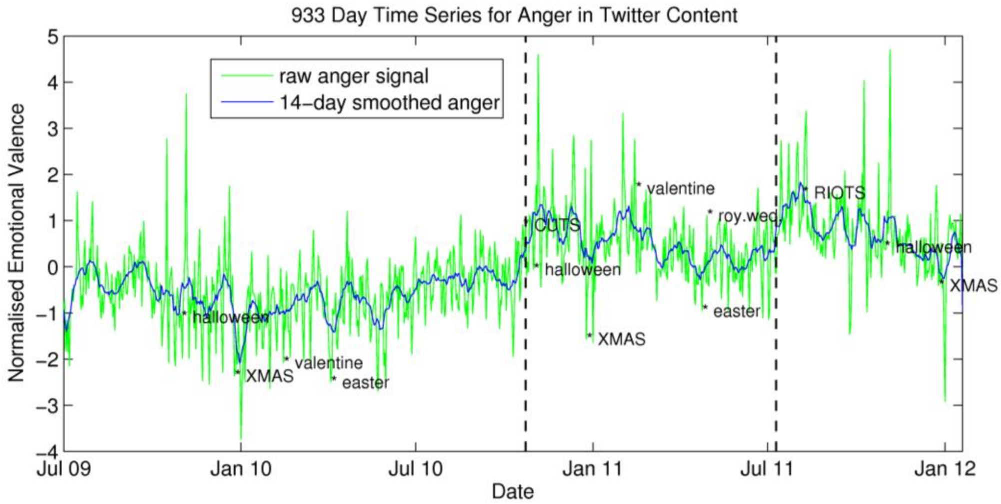


**Figure 3: Plot showing a time series for anger words found in the contents of Twitter in the United Kingdom between 1st July 2009 and 19th January 2012.**

shift in collective mood, particularly fear and anger, are still being felt today, and can be seen in Figure 2 and Figure 3 respectively.

The same testing technique shows a second important change-point, that of August 2011, when riots broke out in various UK cities, leading to looting, arson and even loss of life. While there are a mixture of reasons given for the unrest, ranging from gang culture to the earlier government cuts to public spending (YouGov, 2011), our methods seem to suggest that an increase in public anger preceded - and not followed - the riots. While it is worth noting that this growth in anger seems to have started before the riots themselves, this does not mean that we could have necessarily predicted them. Discovering an interesting correlation after the fact can be of great help to social scientists and other scholars, when interpreting those events, but is very different from predicting the events beforehand. For instance, there have been other increases in anger before, without them leading to any riots. As there is no official record of public mood, we need to be contented with finding correlations between trends in the data for each emotion and events in the external world.

Other social events can also be observed in the data. It appears that the occurrence of an event in early May 2011, possibly the royal wedding between Prince William and Kate Middleton, temporarily paused the ill feelings for a short time, before the rise eventually resumed. This trend points towards the wide media coverage given certain events, such as the royal wedding, altering the national public mood, a currently under-studied effect. We can also find peaks of emotions for other external events, such as the death of Amy Winehouse, and of Osama Bin Laden, both events that received widespread media coverage.

## Conclusions

Our goal was to see if the effects of social events can be seen in the contents of Twitter, and to speculate if some of them could even be predicted. The first part of our analysis provides a sanity check, in that it corroborates our assumption that word-counting methods can provide a reasonable approach to sentiment or mood analysis. While this approach is standard in many applications, we felt that a sanity check in the domain of mood detection via Twitter was necessary. By making use of lists of words that are correlated with the sentiments of Joy, Fear, Anger and Sadness, we observe that periodic events such as Christmas, Valentine and Halloween evoke the same response in the population, year after year.

In the second part of our analysis we were able to observe statistically significant change-points in the collective mood of the nation relating to the announcement of the governmental budget cuts in public spending and the riots taking place across the country in 2011. Each of these events had a large impact on the public mood, with increases in anger and fear still being felt today.

We visualised our findings by combining timelines with face expressions with a tool from the Grimace Project (Spindler, 2009). The end result[1], as seen in Figure 4, can be used by the public as well as by researchers to help their understanding and interpretation of hundreds of millions of tweets at a quick glance.

Social and political sciences can now enter a data-driven phase, but this will inevitably require vast amounts of non-traditional data. The exploitation of Big-Data will require the use of multiple tools, from different fields. Data management, data mining, text mining, data visualisation, all seem to be as necessary as the statistical analysis part (Flaounas, et al., 2011; Sudhahar, Lansdall-Welfare, Flaounas and Cristianini, 2012).

Notice however that since we did not choose the parameters of the mood system so as to correlate our score to the some score for the general UK population, we cannot claim that our mood scores were calibrated to compensate for the various and obvious biases we have in the data collection (unlike in our flu study (Lampos, De Bie and Cristianini, 2010)). All that we can claim – at best – is that we have measured the mood of city dwelling twitter users. They tend to be young; they tend to be savvy and techno-literate; they are definitely a biased sample of the UK population, although a large one, since we included posts by more than 9 million individual users.

The idea that political analysis can assess in real time the mood of millions of people is not without potential for concerns. Serious attention to ethical implications should be paid as these technologies evolve.

[1] Facial expressions and mood timelines for 3 years of UK Twitter data available at http://mediapatterns.enm.bris.ac.uk/mood/

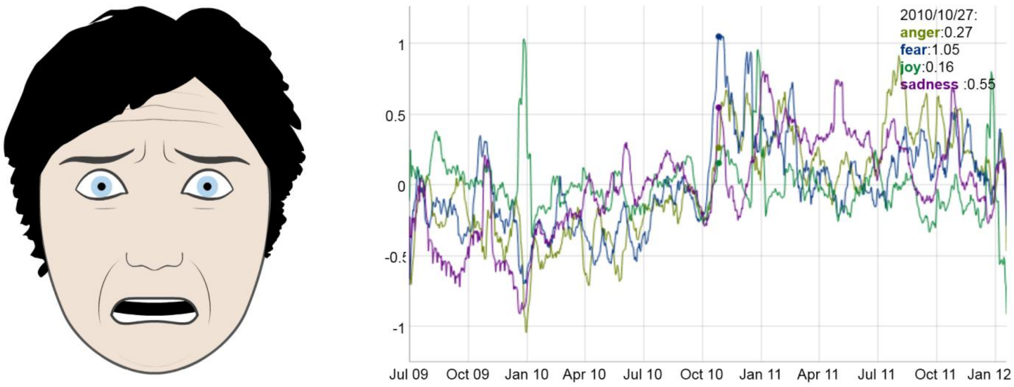


**Figure 4: Screenshot of the facial expression and mood timeline tool available online at [http://mediapatterns.enm.bris.ac.uk/mood](http://mediapatterns.enm.bris.ac.uk/mood).**

## Acknowledgements

N. Cristianini is supported by the CompLACS project (European Community's Seventh Framework Programme -- grant agreement No. 270327); All authors are supported by Pascal2 Network of Excellence.

## References

Bollen, J.; Mao, H.; and Zeng, X.-J. 2010. Twitter mood predicts the stock market. *Journal of Computational Science* 2(1):1–8.

Dodds, P.; Harris, K.; Kloumann, I.; Bliss, C.; and Danforth, C. 2011. Temporal patterns of happiness and information in a global social network: Hedonometrics and Twitter. PloS One, 6(12).

Flaounas, I.; Ali, O.; Turchi, M.; Snowsill, T.; Nicart, F.; De Bie, T.; and Cristianini, N. 2011. NOAM: news outlets analysis and monitoring system. SIGMOD Conference 2011: 1275-1278.

Golder, S. and Macy, M. 2011. Diurnal and seasonal mood vary with work, sleep, and daylength across diverse cultures. *Science*, 333(6051):1878–1881.

Lampos, V.; De Bie, T.; and Cristianini, N. 2010. Flu detector - tracking epidemics on Twitter. In *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases*. Springer, 599–602.

Lansdall-Welfare, T.; Lampos, V.; Cristianini, N. 2012. Effects of the recession on public mood in the UK. Mining Social Network

Dynamics (MSND) session on Social Media Applications in News and Entertainment (SMANE) at WWW '12, pp. 1221-1226, ACM.

O'Connor, B.; Balasubramanyan, R.; Routledge, B. R.; and Smith, N. A. 2010. From tweets to polls: Linking text sentiment to public opinion time series. In *Fourth International AAAI Conference on Weblogs and Social Media*.

Porter, M. 1993. An algorithm for suffix stripping. Program: electronic library and information systems, 14(3):130-137.

Spindler, O. 2009. Affective space interfaces. Unpublished master's thesis, Vienna University of Technology, Vienna, Austria.

Strapparava, C.; and Valitutti, A. 2004. WordNet-Affect: an affective extension of WordNet. Proceedings of the 4th International Conference on Language Resources and Evaluation, pages 1413–1418.

Sudhahar, S.; Lansdall-Welfare, T.; Flaounas, I.; and Cristianini, N. 2012. ElectionWatch: Dectecting patterns in news coverage of US elections. In Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics (EACL), 23-27 April.

YouGov, 2011. The Sun survey results. [online]. Available at : http://cdn.yougov.com/today_uk_import/yg-archives-pol-sun-riots-100811.pdf [Accessed 9 August 2012]