Approved for Public Release; Distribution Unlimited. 12-4711

Early Warning Signals of Tipping-Points in Blog Posts

Dr. David Slater

The MITRE Corporation

Abstract. Complex systems in a variety of disciplines are punctuated by tipping-points, where the system undergoes an abrupt rapid change of state. Early warning signals of such tipping-points can be detected via a concept called critical slowing down by examining a set of statistical metrics. A method based on critical slowing down is used to examine the sentiment of a localized set of blog posts in the lead up to two tipping-points. Evidence of critical slowing down is found in the days leading up to each tipping-point.

1 Introduction

Many complex systems undergo abrupt, rapid transition when they reach a tipping-point (also called a critical transition). Such tipping-points occur in a variety of fields such as Ecology[1, 2], Climatology[3], Finance[4], and Medicine[5]. Prediction of these events is difficult for a number of reasons. First, the system may show little qualitative change prior to the transition[6]. Second, available models of such complex systems are usually not accurate enough to be used for prediction. Third, often only partial (and often peripheral) information is available about the system.

Despite the difficulty in predicting tipping-points, it has been shown that, in many cases, systems have the same generic properties near tipping-points. Scheffer et. al[6] provides a review of the applicable theory as well as a few applications, while Kuehn[7, 8] provides a more mathematical approach by framing the problem in terms of fast-slow dynamical systems. The fundamental idea is that a phenomena called Critical Slowing Down (CSD) may be used to develop early warning signs of tipping-points. As a system approaches a tipping-point, the system slows down and its memory of past events increases. This slowing down can be captured by a set of relatively simple statistical metrics[6, 9]. The use of these CSD metrics to predict slowing down is still in it's infancy. We follow the method of Drake and Griffin[10], where a set of statistical properties are combined into a summary statistic that indicates potential areas of critical slowing down.

We explore CSD as an early warning signal of tipping-points in blog sentiment data. Social media sentiment can be used as a sample of a population's happiness, anxiety, and tension[11–13]. Sentiment is an ideal realm for CSD style analysis as tipping-points in social systems such as protests and rebellion are often preceded

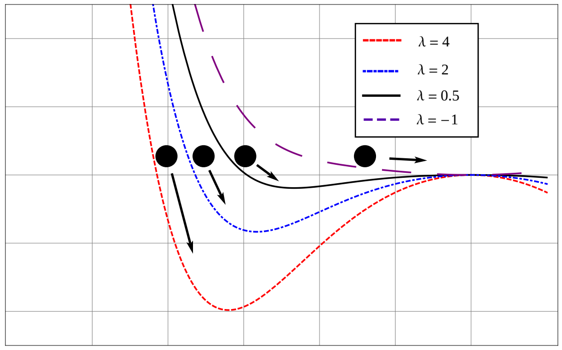

Fig. 1. A system as it approaches a tipping-point ($\lambda = 0$). As λ approaches 0, the well narrows and flattens, causing the ball to take longer and longer to reach the bottom of the well. This phenomena is called Critical Slowing Down.

by a rise of anger, sadness or tension. For this work, we focus on a localized set of blogs and look for early warning signals of tipping-points such as protests and uprising. Predicting such events is of course difficult, but provides a good test case for the use of CSD methods in complex social systems, as such events are often preceded by a period of rising anger, sadness or other emotions. We examine the sentiment of blog posts over a 700 day period, during which two major tipping-points occurred. We find that for both signals, the lead up to each tipping-point was characterized by a period of slowing down. Note that due to data sensitivity, the exact dates and location are not described.

2 Critical Slowing Down Overview

In many cases, when a complex system approaches a tipping-point, its recovery from small perturbations becomes increasingly slow. To illustrate this phenomena, consider a ball rolling on the surfaces in Figure 1 (left). In this example, the bifurcation parameter λ controls the shape of the surface. For $\lambda = 4$, the ball rolls quickly to the bottom of the well, which is both wide and deep. In contrast, as λ approaches 0, the well flattens and narrows, causing the ball to roll slowly to the bottom. When λ passes 0 (the tipping-point), the well disappears and the ball rolls off to the right. This slowing as λ approaches the tipping-point is known as Critical Slowing Down (CSD).

In a complex system, it is impractical to try to measure how quickly the system returns to its steady state as one only has observations about the current and past states. Instead, it has been noted that an increase in a number of statistical metrics occurs near tipping-points[6]. For this work, we use the coefficient of variation, 1-lag autocorrelation, and skewness. An increase in any of these metrics has been observed near tipping-points. Furthermore, we follow

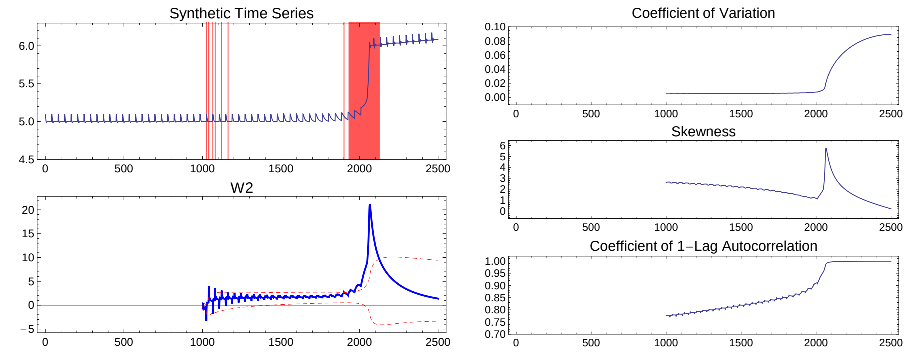

Fig. 2. An example of the CSD method for a synthetic time series. As the series approaches the tipping-point, the three statistical indicators change until the W2 summary statistic detects the system slowing down. This detection is when the blue curve crosses one of the two dashed red curves in the W2 plot (e.g. when the W2 statistic is more than two standard deviations from its long term average).

the method introduced by Drake and Griffin and compute a composite CSD indicator statistic W2 using the method described as follows[10].

First, choose moving window size w, noting that any transitions in the first w points will not be detected. Then, for each window, calculate the coefficient of variation, the lag-1 autocorrelation, and the skewness. The $W2$ statistic is then calculated for all but the first window by computing the sum of the standardized deviations of the current window and from the mean of all previous windows for each measure. Finally, potential critical slowing down areas are identified by finding points where $W2$ is more than two standard deviations from its running average.

2.1 Critical Slowing Down Example

We now return to the case of the ball rolling on a set of surfaces. Figure 2 shows the results of the method on a sample time series. For this time series, the ball is set on the surface corresponding to $\lambda = 1$ and allowed to roll for a fixed period of time. Then λ is decremented by 0.05 and the process is repeated. This continues until λ crosses through 0 (the tipping-point) at time 2100 and it rolls into a different well.

The figure shows the time series, the three metrics computed with a window size of 1000 points, as well as the composite $W2$ statistic. For $W2$, the dashed red lines correspond to two standard deviations from the running average. When W2 is outside two standard deviations of its long run average, a red line is drawn on the time series to indicate a region of possible slowing down; in this case, the tipping-point is detected four generations early.

There are a few issues to note in the figure. First, the method incorrectly identifying CSD several times earlier in the time series. These are false positives

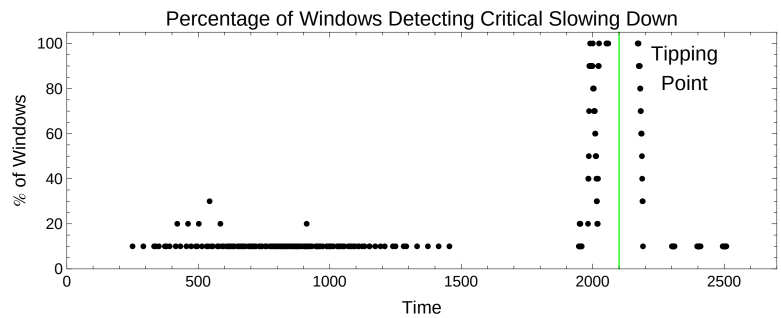

Fig. 3. Results of applying the CSD method for a variety of moving window sizes to the synthetic signal in Figure 1. The CSD metrics are calculated for twelve window sizes ranging from 200 to 1200 points and the percentage of window sizes detecting CSD for each point are shown. The synthetic tipping point is located at time 2100.

due to the standardized deviations being sensitive to a small initial sample set. Second, the optimal window size depends on characteristics of the time series and how best to choose the window size is an open question. Thus, in order to minimize false positives and to account for the uncertainty in choosing the correct window size, we compute the statistics for a variety of window sizes and monitor the percentage of window sizes that detect slowing down for each point. For this example, we vary the window size from 200 to 1200 points in steps of 100. Figure 3 shows the percentage of windows detecting slowing down for each point. Here we see that as we approach the tipping-point at point 2100, nearly all window sizes detect slowing down. We also see that for points between 200 and 1500, few window sizes detect slowing down suggesting these are false positives. By using a multitude of window sizes we more accurately predict where CSD may be occurring.

3 Blog Analysis

The CSD method was used on a model of blog sentiment, obtained through linguistic analysis, social network analysis, ideology and topic detection, sentiment classification and metadata extraction. A localized set of blog posts were analyzed for sentiment words, from which daily positive and negative sentiment scores were computed. A total of 11,417 blog posts from 102 unique blogs spanning 700 days were analyzed for their sentiment content. Positive and negative daily sentiment scores were computed as the percentage of the total words that had either positive or negative sentiment. Figure 4 shows histograms for the daily sentiment scores. Both the positive and negative scores are normally distributed, with means of 12.55 and 8.28 respectively. Interestingly the percentage of positive and negative words are also positively correlated with a correlation coefficient of 0.44.

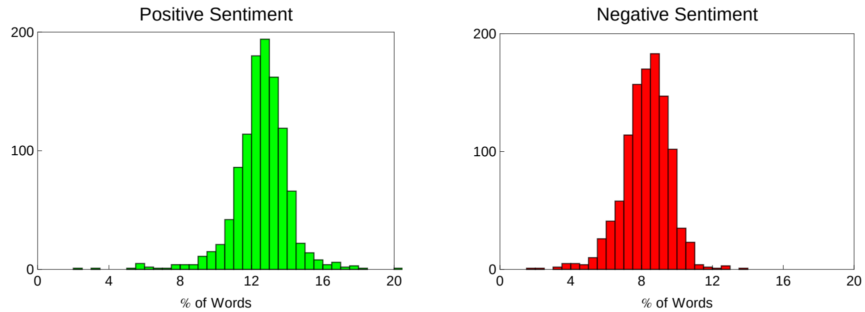

Fig. 4. Distribution of daily positive (left) and negative (right) sentiment scores over the 700 days considered.

Figure 5 shows the number of blog posts as well as the positive and negative sentiment scores for the 700 days considered. Over this time period, subject matter experts identified two major tipping-points on days 529 and 696. It is important to note that examination of the sentiment time series gives little indication of a potential tipping-point as they are not punctuated by spikes in sentiment. Note that two other significant events occurred on days 447 and 642, punctuated by a significant spike in posts. These are non tipping-point induced spikes (e.g. death of an important figure) and not of interest in this work.

We apply the CSD method to the number of blog posts, positive sentiment score, and negative sentiment score time series for the months leading up to each of these tipping-points. Before applying the CSD method, we used a robust local regression smoother on the time series, smoothing the data with a span of 3 days and ignoring outliers that are 6 standard deviations from the mean. This is done for two reasons. First, there is some lag inherent in blog postings. Second, outliers such as those that occur on days 447 and 642 decrease the effectiveness of the method since as those days move in and out of the moving window, they have a disproportionate effect on the the metrics.

3.1 Tipping-Point 1: Day 529

The CSD method is first applied to the lead up to the event on day 529. As with the example in section 2, we calculate the W2 statistic for a variety of window sizes. For each day the percentage of windows that detect critical slowing down is then determined. Figure 6 shows the results for the lead up to the tipping-point on day 529, where the moving window size is varied from 30 days to 250 days in steps of 10 days. We find that for the negative sentiment score that on day 494, 33 days before the tipping-point a significant number of windows have detected the statistical properties changing, signaling the system may be slowing down as it approaches a tipping-point. The detection of critical slowing down is found for all days up through the event on day 494.

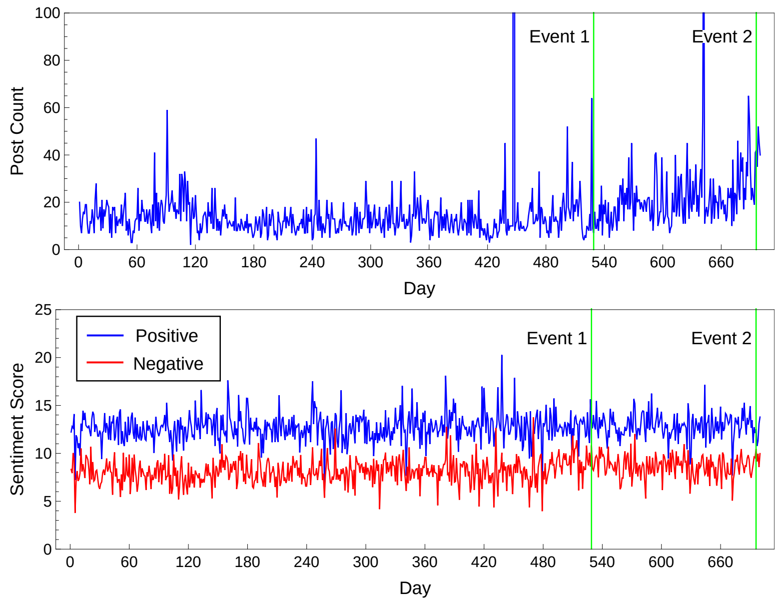

Fig. 5. Time series of the number of blog posts per day as well as the daily positive and negative sentiment scores. The green lines correspond to tipping-points on days 529 and 696.

For the positive sentiment scores and the daily post count time series, the CSD method does not detect early warning signals of the tipping-point. The positive score series does have non trivial peaks around days 180 and 250, which coincide with other less significant events. The blog post time series does not significantly detect critical slowing down at any time before the tipping-point.

It is again important to note that for each series, it is expected that for most days, a few windows may detect slowing down. This is due to the changing window size moves where W2 statistic begins being calculated and the fact that the W2 statistic typically moves outside 2 standard deviations from its mean in the first few iterations. These false positives could be excluded by excluding the first 10% of points from being a potential slowing down region.

3.2 Tipping-Point 2: Day 696

CSD analysis was also conducted for the lead up to the tipping-point on day 696. Here we take as our time series days 540 through 700 so that we exclude the first tipping-point. The percentage of windows detecting critical slowing down for the three time series is shown in Figure 7, where the moving window size is varied from 30 to 130 in steps of 10. As with the first tipping-point, we find that the negative sentiment signal detects critical slowing down twelve days prior to the tipping-point. In addition, 38 days before the tipping point the post count

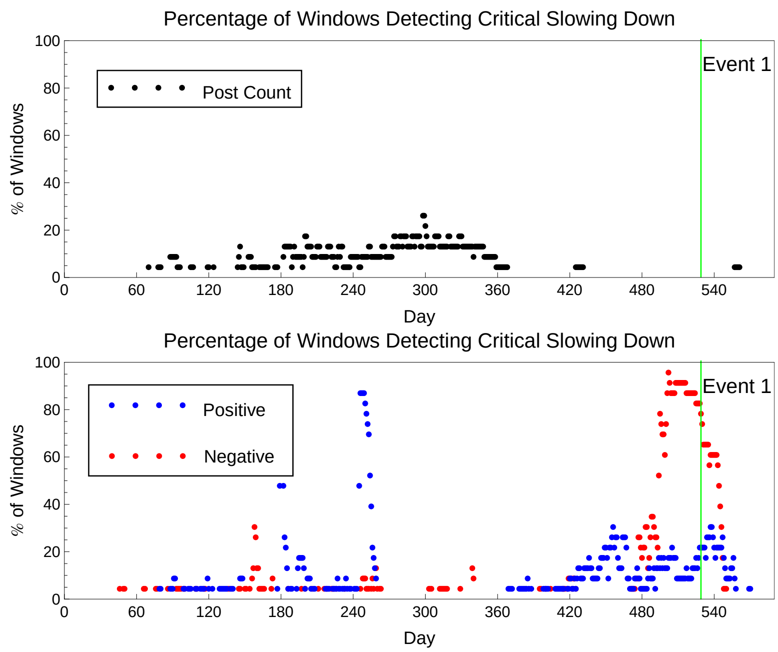

Fig. 6. Results of applying the CSD method to the daily post count, negative sentiment, and positive sentiment score for the lead up to the tipping-point on day 529. For each, the moving window size was varied from 30 to 250 days in steps of 10 days. The percentage of windows detecting slowing down for each day is displayed, showing the negative sentiment time series detects the tipping-point 33 days early.

series also detects evidence of critical slowing down, suggesting that this is also a useful metric. Finally, the positive sentiment signal has a spike at day 608.

4 Concluding Remarks

The concept of critical slowing down has been used in a variety settings to detect early warning signals of tipping-points and other critical transitions. We apply critical slowing down to a new domain (social systems) by exploring the detection of tipping-points in the sentiment of a localized set of blog posts. The results are promising as the CSD method detects evidence of critical slowing down prior to the two tipping-points considered.

Further investigation is needed to make critical slowing down based methods robust. In most of the current applications of critical slowing down, a tipping-point is identified and lead up to the tipping-point is analyzed. In order to move to predictive analysis, a better understanding of the robustness of CSD metrics in situations where spatial complexity, chaos, and stochastic perturbations govern the dynamics is needed.

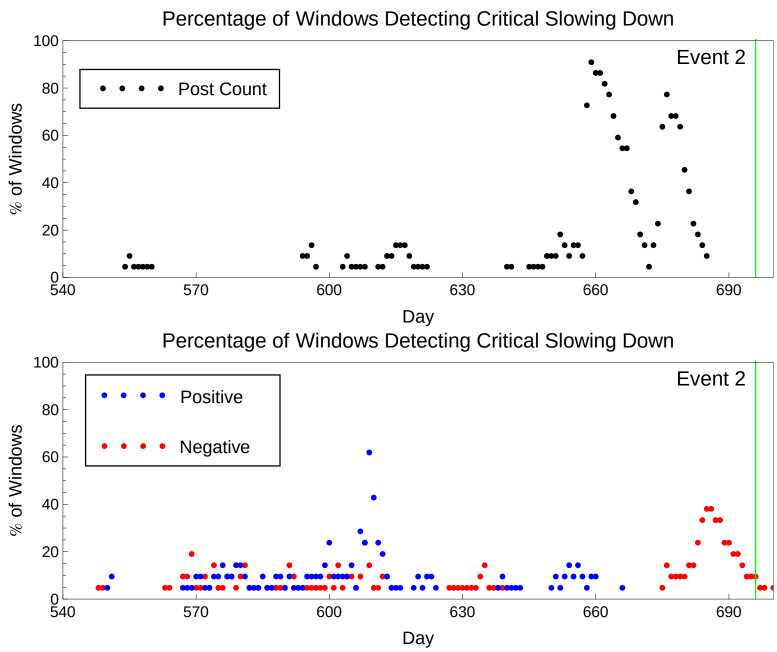

Fig. 7. Results of applying the CSD method to the daily post count, negative sentiment, and positive sentiment score to the lead up to the tipping-point on day 696. For each, the moving window size was varied from 30 to 130 days in steps of 10 days. The percentage of windows detecting slowing down for each day is displayed, showing the daily post count time series detects the tipping-point 17 days early.

The robustness of the W2 statistic in particular is still an open question. Statistical procedures for reducing false positives and negatives need to be developed and the relationship between window size and signal properties needs to be established. Finally, how best to deal with signals with multiple tipping-points is yet to be determined.

References

1. Scheffer, M., Carpenter, S., Foley, J.A., Folke, C., Walker, B.: Catastrophic shifts in ecosystems. Nature **413**(6856) (October 2001) 591–6
2. Scheffer, M., Carpenter, S.R.: Catastrophic regime shifts in ecosystems: linking theory to observation. Trends in Ecology & Evolution **18**(12) (December 2003) 648–656
3. Lenton, T.M., Held, H., Kriegler, E., Hall, J.W., Lucht, W., Rahmstorf, S., Schellnhuber, H.J.: Tipping elements in the Earth's climate system. Proceedings of the National Academy of Sciences of the United States of America **105**(6) (February 2008) 1786–93

4. May, R.M., Levin, S.A., Sugihara, G.: Complex systems: Ecology for bankers. **451**(7181) (February 2008) 893–895
5. McSharry, P.E., Smith, L.a., Tarassenko, L.: Prediction of epileptic seizures: are nonlinear methods relevant? Nature medicine **9**(3) (March 2003) 241–2; author reply 242
6. Scheffer, M., Bascompte, J., Brock, W.A., Brovkin, V., Carpenter, S.R., Dakos, V., Held, H., van Nes, E.H., Rietkerk, M., Sugihara, G.: Early-warning signals for critical transitions. Nature **461**(7260) (September 2009) 53–9
7. Kuehn, C.: A mathematical framework for critical transitions: Bifurcations, fast-slow systems and stochastic dynamics. Physica D: Nonlinear Phenomena **240**(12) (June 2011) 1020–1035
8. Kuehn, C.: A mathematical framework for critical transitions: normal forms, variance and applications. ArXiv e-prints (2012)
9. Scheffer, M.: Complex systems: foreseeing tipping points. Nature (2010) 6–7
10. Drake, J.M., Griffen, B.D.: Early warning signals of extinction in deteriorating environments. Nature **467**(7314) (September 2010) 456–9
11. Balog, K., Mishne, G., de Rijke, M.: Why are they excited?: identifying and explaining spikes in blog mood levels. ... of the Eleventh Conference of the ... (2006) 207–210
12. Dodds, P.S., Danforth, C.M.: Measuring the Happiness of Large-Scale Written Expression: Songs, Blogs, and Presidents. Journal of Happiness Studies **11**(4) (July 2009) 441–456
13. Dodds, P.S., Harris, K.D., Kloumann, I.M., Bliss, C.a., Danforth, C.M.: Temporal patterns of happiness and information in a global social network: hedonometrics and Twitter. PloS one **6**(12) (January 2011) e26752