

# Carbon: Forecasting Civil Unrest Events by Monitoring News and Social Media

Wei Kang<sup>1</sup>(✉), Jie Chen<sup>1</sup>, Jiuyong Li<sup>1</sup>, Jixue Liu<sup>1</sup>, Lin Liu<sup>1</sup>, Grant Osborne<sup>2</sup>,  
Nick Lothian<sup>2</sup>, Brenton Cooper<sup>2</sup>, Terry Moschou<sup>2</sup>, and Grant Neale<sup>2</sup>

<sup>1</sup> University of South Australia, Mawson Lakes, SA 5095, Australia  
{wei.kang,jie.chen,jiuyong.li,jixue.liu,lin.liu}@unisa.edu.au

<sup>2</sup> Data to Decisions CRC, Kent Town, SA 5067, Australia  
{grant.osborne,nick.lothian,brenton.cooper,terry.moschou,  
grant.neale}@d2dcrc.com.au

**Abstract.** Societal security has been receiving unprecedented attention over the past decade because of the ubiquity of online public data sources. Much research effort has been taken to detect relevant societal issues. However, forecasting them is more challenging but greatly beneficial to the entire society. In this paper, we present a forecasting system named Carbon to predict civil unrest events, e.g., protests and strikes. Two predictive models are implemented and scheduled to make predictions periodically. One model forecasts through the analysis of historical civil unrest events reported by news portals, while the other functions by detecting and integrating early clues from social media contents. With our web UI and visualisation, users can easily explore the predicted events and their spatiotemporal distribution. The demonstration will exemplify that Carbon can greatly benefit the society such that the general public can be alerted in advance to avoid potential dangers and that the authorities can take proactive actions to alleviate tensions and reduce possible damage to the society.

**Keywords:** Civil unrest · Predictive models · Open source data

## 1 Introduction

Societal security has great impact on daily lives of the public and the stability of a society. Great attention has been paid to security incidents thanks to instant news reports and the widespread use of social media. Nowadays, people can read “hot-off-the-press” emergency reports instantly via online news portals, while social media users can even smell the emergence of incidents earlier before their occurrences. Much effort has been invested in detecting events or generating summaries from social media data [4–6,8]. Compared with detecting, however, forecasting them could be much more desirable, as it allows the public to be alerted prior to dangers and enables the authorities to take proactive actions to alleviate tensions and minimise disruption. Some researchers tried to make forecasts with public data to predict crowd behavior [3]. Researchers from Virginia

Tech has built EMBERS to forecast civil unrest events such as protests in 10 Latin American countries [1, 7]. As an intelligence project supported by IARPA<sup>1</sup>, other than the published papers, however, plenty of details of EMBERS are still unclear and kept confidential, and the system is not accessible by external users. Besides, it focuses on forecasting only for Latin American countries.

In this paper, we have built a system named Carbon to forecast civil unrest events for Australia and some other Asia-Pacific countries by monitoring and analysing news and social media. To process millions of civil unrest related documents each day, Carbon has been built on Apache Spark<sup>2</sup> to provide real-time and scalable data processing, analysing and event forecasting. Note that civil unrest forecasting is more challenging for Asia-Pacific countries, esp. Australia, where such events are often less frequent and on smaller scales. To tackle the challenges, we propose a novel and effective time-series model and improve the planed protest model in EMBERS for fine-grained predictions.

Carbon ingests open source data, mainly news articles and Twitter/Facebook streams, to capture precursory clues for civil unrest events. Our news analysts search through popular news portals every day and take down reported civil unrest events as Gold Standard Records (GSRs), which are then used as the ground truth to build and evaluate our models. Each GSR represents a reported event, e.g., a protest, with attributes such as the event date, reported date, location, predefined event type (e.g., political or economic issues), predefined population group (e.g., education, labor) and description. We also purchase Twitter data from GNIP<sup>3</sup> continuously to avoid the Twitter API rate limits and to collect as much civil unrest related data as possible to enhance the predictiveness of our system. Carbon is currently aimed at Australia and nearby countries. However, with our generic system design and solutions, it can be easily extended to other regions.

To predict civil unrest events, we propose a GSR based model which utilises underlying patterns of historical GSRs, and another model which leverages civil unrest indicators extracted from social media streams. These models are designed to work complementarily. That is, the GSR based model makes use of evolutionary features of different types of historical events, and the other model analyses social media precursory clues. We notice that, although some events can be predicted by both models, most of them are captured by one model only, indicating that the two models can complement each other to provide a greater coverage.

Since predictions are made to forecast future events, the evaluation of the predictions cannot be conducted until GSRs are collected for the corresponding predicted time period. A background job is scheduled weekly in Carbon to evaluate the precision and recall of those predictions. A prediction is considered correct only if all its attributes (mainly the date or time range, location, event type and population group) match the corresponding attributes of some GSR within the examined time period. In addition, we compute a lead time for

<sup>1</sup> <https://www.iarpa.gov/>.

<sup>2</sup> <http://spark.apache.org/>.

<sup>3</sup> <https://gnip.com/>.

each true prediction, i.e., how many days in advance the prediction is produced before the corresponding event is reported on news portals. The quantitative experimental results are summarised in Table 1.

The targeted audience includes both the public and authorities, who will have a deeper understanding of the importance of a civil unrest forecasting system to the maintenance of the security and stability of a society. We will demonstrate how to explore the predictions through a web UI with three views, i.e., the Predictions List showing recent predictions with details such as model name, generation/predicted event date, location, probability, supporting evidence, etc., and the Predictions Map and Timeline which visualise the geographical and temporal distribution of the predictions respectively.

Our major contributions include: (1) we have built a system which forecasts rather than detects civil unrest events for countries where civil unrest events are less frequent; (2) we have introduced two predictive models which make use of different types of knowledge (i.e., evolutionary trends and social media precursory clues of civil unrest events respectively) and work complementarily; (3) the predictions are of great value for the benefit of both the general public and the authorities.

## 2 System Architecture

As is shown in Fig. 1, the architecture of Carbon consists of three major components – the data storage, the data processing and modelling, and the Carbon UI. The data processing and modelling component in the middle is the core of Carbon. Spark jobs are scheduled and submitted periodically to perform day-to-day operations, such as continual ingestion of open source data, data enrichment and feature extraction (including the inference of Twitter/Facebook users’ locations based on their profiles and check-ins, the conversion of a relative day, e.g., next Monday, to an absolute date, etc.), inverted index building and so on.

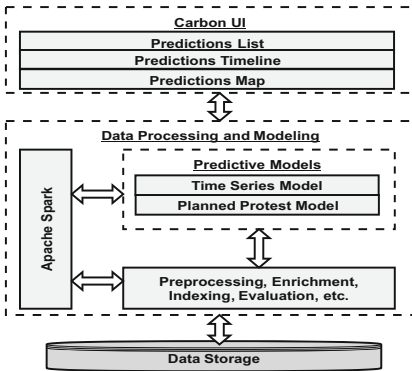


Fig. 1. Architecture of Carbon

Table 1. Experimental results

| Model                   | Precision | Recall | Lead time |
|-------------------------|-----------|--------|-----------|
| Time-series or TS (0/1) | 0.42      | 0.27   | 5         |
| TS (approx.)            | 0.42      | 0.55   | 5         |
| TS (1 month)            | 0.8       | 0.52   | 5         |
| PPM                     | 0.38      | 0.94   | 9         |

The predictive models which are designed to work complementarily are submitted to run at regular intervals by a scheduler. The output predictions are stored in the database and presented in the web UI, and will be scheduled later for evaluation when GSRs are gleaned for the predicted time period.

### 3 Predictive Models

#### 3.1 Time-Series Model

The time-series model predicts the future by leveraging both historical evolutionary patterns and recent trends of events. It first sorts the GSR events ranging from  $t_{start}$  to  $t_{end}$  in chronological order, and divides them into equally sized (e.g., one week long) windows  $w_1, w_2, \dots, w_n$  such that each window contains events falling in the corresponding time range. With the  $n$  windows, our aim is to predict whether certain events will happen in window  $w_{n+1}$  by exploring historical events based on a group of base patterns. A base pattern is produced by examining GSR events in a few, say 3, recent windows to capture the latest trends of a certain type of events. For instance, if a series of protests was about immigration officers calling for higher salary and it happened in windows  $w_{n-2}$  and  $w_n$  but not  $w_{n-1}$ , we can derive a time-series base pattern  $(1, 0, 1)$ , where 1 indicates the occurrence of a protest in corresponding windows and 0 otherwise, to capture recent trend of this type of protests. To make predictions, we need to generate all possible base patterns in recent windows. Once having the base patterns ready, we then scan GSR events window by window to figure out the number of occurrences of the same patterns in history. For each base pattern  $bp$ , we obtain the numbers of cases that the same type of events happened or did not happen respectively in the subsequent window when there are consecutive windows matching  $bp$  in history, and then estimate the probability of the corresponding event to happen in window  $w_{n+1}$ . We have adopted the idea of higher order Markov chains in this model, meaning the prediction of  $w_{n+1}$  depends on multiple precedent windows  $w_n, w_{n-1}, \dots$  instead of only the most recent window  $w_n$  because the occurrence of an event in the future is likely to be triggered by similar events that happened in multiple previous windows.

To enhance the sensitivity of the model, we further extend the 0/1 base pattern matching approach to an approximate matching approach. We represent a base pattern using the frequency of a certain type of events,  $(1, 0, 4)$  for instance, to match against the frequencies of the  $n$  windows of GSR events of the same type. We support approximate frequency matching by allowing base pattern  $(1, 0, 4)$  to be matched with some sub-history (i.e., a subsequence of the GSR history), say  $(\dots, 2, 0, 5, \dots)$ , as long as the sum of corresponding absolute frequency differences, i.e.,  $|2 - 1| + |0 - 0| + |5 - 4|$ , is less than a distance threshold. Besides, we also support approximate matching in the timeline dimension by allowing  $(1, 0, 4)$  to be matched with some sub-history of a different length, say  $(\dots, 2, 0, 0, 5, \dots)$ . A combined difference is computed and compared with the threshold to take into account the differences in both the frequency and timeline dimensions, which provides a fine-grained comparison between a base pattern

and sub-histories by matching their signal strengths and durations with certain tolerance.

As shown in Table 1, the experimental study has verified that the approximate matching approach can significantly improve the average recall. As a trade-off, if we gradually expand the prediction time window from one week to one month, the average precision will continuously increase to 0.8 while the average recall slightly drops to 0.52.

### 3.2 Two-Phase Planned Protest Model

Although most types of events occur from time to time, there are still certain types which never appeared before or cannot be modelled by the GSR based model. As such, we have also introduced a *two-phase* planned protest model (PPM) which clusters social media textual data to generate predictions with strong evidence. Our PPM is an extension of the planned protest model in [7] to overcome the difficulty in distinguishing different predictions generated for the same location and event date.

PPM applies two clustering phases over the indicators, which are the enriched data containing the original contents (e.g., Tweets) and inferred phrase lists. A list often contains seed keywords and related lemmas (e.g., “good” is the lemma of “better”), inferred future dates and locations. The indicators are filtered to English only, and binned into date and location groups. In the first phase, we extract cleaned (i.e., with unicode symbols, RT blocks, quotes, etc. removed) textual Tweet features first and maintain them using a token/count vectoriser. Next, we estimate an epsilon density measure using a nearest neighbour max absolute distance and apply it in DBSCAN [2] to cluster Tweet based indicators whose textual contents are similar. In the second phase, we add new data sources (e.g., Facebook URLs) to existing Tweet clusters based on the URLs contained in the Tweets, and merge related Tweet clusters together if they all reference the same URL (e.g., a news article). In the end, these clustered indicators are output as predictions if they cross over a configurable threshold. The performance of PPM can be found in the last row of Table 1.

## 4 Demonstration

In the demonstration, we will exhibit Carbon, our civil unrest forecasting system, which has been producing predictions since June 2016. We will elaborate the system design, predictive models and predictions, such that users can obtain an in-depth understanding of how the system effectively generates predictions and why the predictions it produces can greatly benefit the society.

We will also introduce the system interface, where we provide user admin functions to grant proper privileges to a new user. Users can log in to see the main interface (cf. external snapshots<sup>4</sup>), where on the left panel users can click

<sup>4</sup> <https://www.dropbox.com/s/7pwjnj97r46abmo/carbonUI.pdf?dl=0>.



(a) Prediction issued by Time-series Model

(b) Prediction issued by PPM

Fig. 2. Example predictions

on one of the three views, i.e., the Predictions List, Map and Timeline, while on the right details of the selected view will be presented. At the top of each view is a section where users can specify a time range and choose one or both of the models to explore the predicted events.

The Predictions List displays the detailed predictions which are likely to happen in the coming few weeks in chronological order. As an example, Fig. 2 shows two predictions, issued by the time-series and PPM models respectively, regarding the same bus drivers strike planned for 4 April 2017 in Adelaide. It was reported on 1 April that up to 50,000 commuters would face great chaos after bus drivers voted for a 24-hour strike, which could halt a third of Adelaide Metro bus services<sup>5</sup>. The time-series model issued the prediction on 29 March, 6 days before the planned strike, based on 30 historical GSRs, and predicted “labour” as its population group which referred to the bus drivers. Meanwhile, the PPM model also predicted the same event on 2 April based on 5 Tweets, from which keywords were extracted and visualised in word cloud. Although the strike was called off owing to an in-principle agreement between the bus service operator and the workers union right before it was about to take place, similar events are not always avoidable. Had the strike happened, great chaos could hit the city, let alone the possibility that it might have triggered conflicts and other violent events. Therefore, forecasting strikes and other civil unrest events can be vital and greatly beneficial to the public and authorities for them to take proactive actions to avoid possible loss and disruption.

Different from the List view, the Predictions Map and Timeline views integrate multiple predictions and visualise their spatiotemporal distributions over a map and a histogram timeline. By hovering the mouse over a circle on the map or a bar in the timeline, users will see integrated information about a

<sup>5</sup> <http://www.adelaidenow.com.au/news/south-australia/09a3a63b361a785fab949828148025e>.

location/date, including the number of predictions, their corresponding models and probabilities. Users can gain insight into safety condition of different places by observing the distribution of reported and predicted civil unrest events. The observations could serve as guidance for travel planning.

## References

1. Doyle, A., Katz, G., Summers, K., Ackermann, C., Zavorin, I., Lim, Z., Muthiah, S., Zhao, L., Lu, C.T., Butler, P., et al.: The EMBERS architecture for streaming predictive analytics. In: *IEEE BigData*, pp. 11–13 (2014)
2. Ester, M., Kriegel, H., Sander, J., Xu, X.: A density-based algorithm for discovering clusters in large spatial databases with noise. In: *SIGKDD*, pp. 226–231 (1996)
3. Kallus, N.: Predicting crowd behavior with big public data. In: *WWW*, pp. 625–630 (2014)
4. Kang, W., Tung, A.K., Chen, W., Li, X., Song, Q., Zhang, C., Zhao, F., Zhou, X.: Trendsmedia: an internet observatory for analyzing and visualizing the evolving web. In: *ICDE*, pp. 1206–1209 (2014)
5. Kang, W., Tung, A.K., Zhao, F., Li, X.: Interactive hierarchical tag clouds for summarizing spatiotemporal social contents. In: *ICDE*, pp. 868–879 (2014)
6. Nguyen, D.T., Jung, J.J.: Real-time event detection on social data stream. *Mob. Netw. Appl.* **20**(4), 475–486 (2015)
7. Ramakrishnan, N., Butler, P., Muthiah, S., Self, N., Khandpur, R., Saraf, P., Wang, W., Cadena, J., Vullikanti, A., Korkmaz, G., et al.: ‘Beating the news’ with EMBERS: forecasting civil unrest using open source indicators. In: *SIGKDD*, pp. 1799–1808 (2014)
8. Zhou, X., Chen, L.: Event detection over twitter social media streams. *VLDB J.* **23**(3), 381–400 (2014)