Increasing Lead Time and Granularity of Civil Unrest Prediction through Time Series Data

Lu Meng
State University of New York at Buffalo
Buffalo, New York
lumeng@buffalo.edu

Rohini K. Srihari
State University of New York at Buffalo
Buffalo, New York
rohini@buffalo.edu

ABSTRACT

Surprise events such as Ferguson Uprising, October 2000 events in Israel, and Iraq's invasion of Kuwait, can have critical impact on the society and people's daily lives. Existing civil unrest prediction systems have so far relied on manually curated data sources and heuristic features determined by human expertise. Current research focuses on machine learning approaches which do not effectively use heterogeneous data sources reflecting varying time series data, and only pays little attention to unstructured text data. In this paper, we propose: (a) a novel approach to model structured and unstructured data, and (b) a predictive model which effectively exploits such data and predicts with increased lead time through LSTM networks. To the best of our knowledge, none of the existing work focused on both data types in the area of civil unrest prediction. Extensive experiments have been conducted on 2 different datasets for both country-level and city-level to illustrate the effectiveness of our model. Code and data is available on GitHub [1].

CCS CONCEPTS

• **Applied computing** → **Sociology**; **Forecasting**.

KEYWORDS

Civil unrest events prediction, social media, multiple data sources, recurrent neural network, deep learning, word embeddings

ACM Reference Format:

Lu Meng and Rohini K. Srihari. 2019. Increasing Lead Time and Granularity of Civil Unrest Prediction through Time Series Data. In *MileTS '19: 5th KDD Workshop on Mining and Learning from Time Series, August 5th, 2019, Anchorage, Alaska, USA.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/1122445.1122456

1 INTRODUCTION

Civil Unrest (CU) events are physical acts that occur in public venues such as demonstrations and protests. Having the ability to provide predictions and early warnings about such events has significant impact in both the public and private sectors. It also benefits NGOs, relief agents in allocating resources to conflict areas.

CU events are believed to be the comprehensive results of multiple factors, which could be related to economic, environment, etc. Widespread of social media has made it as another important factor as it provides people a platform to orchestrate, announce, and organize activities. The heterogeneous nature and complex mechanism of CU events make it challenging to develop the predictive systems. In recognition of the importance of this task, IARPA held a competition, the Mercury Challenge[2], which encouraged innovative solutions for the automation of societal events forecasts in the Middle East.

In this paper, we are trying to answer the following questions:

- Is there any pattern we can learn to forecast CU events from the data in our day-to-day life, such as general social sentiment, economic, and politics?
- How can we represent such data?
- Given a location (i.e. country, city), can we predict the daily CU event counts accurately in advance?

The main contributions of our study are summarized as follows: 1) We propose a novel approach to model both structured data and unstructured text data. 2) We study proxies of various societal aspects with different update rates by introducing the deep learning model, which learns the representations and features of time series data automatically, in contrast of the traditional ways which rely heavily on human expertise and efforts. 3) In participation of the IARPA Mercury Challenge, we are able to forecast the event counts for a specific day/week over a country/city, which poses more challenges and significance than only giving whether there will be an event in such location.

2 RELATED WORK

Studies have been done for the purpose of predicting offline civil unrest events by using various kinds of data sources. Information extracted form the Global Database of Events, Language, and Tone (GDELT) [8] is widely used by research work on predicting protest events, such as the temporal burst pattern of GDELT [11]. Some work focuses on using external information to predict conflicts, for example, future climate scenarios and variables like population are used to forecast conflicts in sub-Saharan African [16]. Significant positive correlation between the volume of future protests descriptions on social media and protest onsets has already been proved [17]. Features such as daily volume or number of posts containing keywords are extracted from Twitter for forecasting social unrest events [20] [6] [19].

[1] https://github.com/conference-submission/submission

[2] https://www.iarpa.gov/challenges/mercury.html

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
MileTS '19, August 5th, 2019, Anchorage, Alaska, USA
© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-9999-9/18/06…$15.00
https://doi.org/10.1145/1122445.1122456

In the work described above, traditional machine learning models are used, such as logistic regression models [17], Hidden Markov model [12], and Random forest classifier [5]. Some studies have applied Dynamic Query Expansion on Twitter data in order to capture emerging conditions for protests by dynamically growing the vocabularies of interest [14] [19] [4].

Gaps in current research: Current approaches pose the limitation in combining heterogenous data sources, and addressing the importance of time sensitive sequential data and text data. Some systems only have the ability of forecasting civil unrest event on country level [2]. Others have conducted similar work but benchmarked on event datasets that are static, and only focus on answering questions such as "will there be a protest in the next few days". This motivates our work presented in this paper to study various potential drivers of civil unrests by using deep learning model which overcomes the shortcomings of traditional systems.

3 DATA SOURCES DESCRIPTION

We consider CU events in 2 countries (Egypt, Jordan) and 3 cities (Cairo, Amman, Delhi). Our work involves more than 1395 data points from over 46 months timespan from May 01, 2015 to February 23, 2019 (ACLED data for Jordan and India starts in January, 2016).

Historical data. We use a large volume of data generated by IARPA for the Mercury Challenge as the ground truth for our evaluation. The Gold Standard Reports (GSR) data sets contain details on more than 120,000 significant events in the areas of Military Activity, Disease, and Non-violent Civil Unrest, which is the type of events we focus on in this paper. Each event has the information about the date, type, location, first reported date, news source, etc.

A second dataset, the Armed Conflict Location & Event Data Project (ACLED) data [13], is included to validate our observations and deploy our framework for city specific predictions. ACLED data is a publicly available dataset which contains date, actors, locations, and fatalities of all reported violence and protest events across Africa, the Middle East, and so on. Comparing to the GSR data, ACLED data has better coverage of countries, and richer description of the data. One problem of using ACLED for forecasting future events is that there is a lag of one week or more for availability of data, which serves as one of our motivations to include dynamic social media data.

Social media data. Social media has transformed the traditional way for people to express their political and social concerns [18]. We collect political tweets from a list of accounts who are politicians and journalists in Egypt and Jordan. All tweets from the 3 cities are also collected.

Open Source Indicators. Open Source Indicators such as Google Trends[3] have the potential to reveal the dynamics of social behavior that precede civil unrest [9]. We gathered Google Trends data in Egypt, Jordan, and India for each keyword set, which will be discussed in detail in 4.2.

[3]Source: https://www.google.com/trends

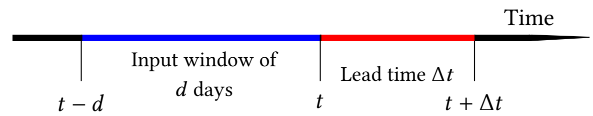

Figure 1: Prediction timeline of CU event counts. t: present time, d: input window size, Δt: lead time.

Economic indicators. Studies show that exogenous political and economic shocks can serve as the necessary underlying drivers of social unrest [3]. Commodity prices[4], unemployment rate, inflation rate[5], etc. are collected for each country.

4 METHODOLOGY

4.1 Problem definition

Our goal is to forecast and generate early warnings of CU event counts by ingesting multiple data sources. We define this task as a classification problem: suppose $\mathbf{x}_t \in \mathbb{R}^n$ is a feature vector corresponding to a day t, where each entry $x_i \in \mathbb{R}^d$ of the vector contains the historical data from the previous d days, and $y_{t+\Delta t} \in \{0, \ldots, k\}$ refers to the category indicating the number of events on the day of $t + \Delta t$. We are trying to produce the function $f : \mathbb{R}^n \to \{0, \ldots, k\}$, such that

$$y_{t+\Delta t} = f(\mathbf{x}_t).$$

An illustration of the prediction timeline is shown as Figure 1.

4.2 Input Features

We classify the data mentioned in Section 3 into the following two categories, structured data and unstructured text data. Features of each data source are extracted as follows and summarized in Table 1.

4.2.1 Historical data. Daily event counts for Egypt and Jordan are aggregated from the GSR dataset. Events that are categorized as "Riots" and "Protests" in ACLED, are extracted for Cairo, Amman, and Delhi.

4.2.2 Keywords selection using Google Trends data. The initial keyword sets denoting civil unrest, are determined by political scientists and CIA The World Factbook[6]. Each set contains the English words, synonyms, and Arabic translations. We then collect weekly Google Trends data for each keyword set per country. Covariance between Google Trends of different keyword sets and GSR event counts are calculated. Keyword sets that are positively correlated are then selected to build the keyword dictionary. For example, "conflict, conflicts, ...", "protest, protester, ..." are selected for Egypt.

4.2.3 Social media data. Daily volume and sentiment (percentage of angry posts) are calculated by filtering and aggregating the political tweets by the keyword dictionary. The unstructured text data, i.e., the relevant tweets, are extracted by filtering all the tweets using the keyword dictionary.

[4]Source: http://www.worldbank.org
[5]Source: https://tradingeconomics.com
[6]Source: https://www.cia.gov/library/publications/resources/the-world-factbook/

Table 1: Feature Taxonomy in Input Data

Category	Date source	Features
Structured	Historical data	CU event counts in GSR and ACLED data.
	Social media data	Daily volume of political tweets and the sentiment.
	Economic indicators	Commodity prices, unemployment rate, etc.
Unstructured	Social media data	Twitter data embeddings of all relevant tweets.

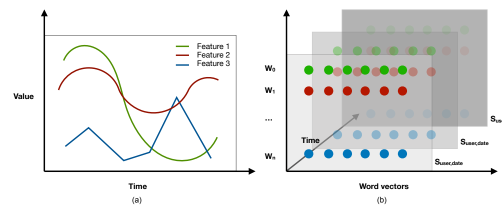

Figure 2: Two ways to model sequential data: (a) structured data, (b) unstructured text data.

4.2.4 Economic indicators. We calculate the covariance between GSR event counts and major economic indicators for each country according to CIA The World Factbook. It can be observed that the relativeness of indicators vary from country to country. Therefore, different features which are positively correlated are selected for different counties.

4.3 Data Vector Representations

There is an analogy between the strokes of composite sketches and the changes of civil unrest indicators. Such as the ordering among the strokes when we draw the sketches and among the occurrences of indicators' change, i.e., the change of inflation rate could lead to the change of unemployment rate. Additionally, contributions to the final sketch differ among different characteristics in eyewitnesses' memory, as well as among the social indicator statistics. For example, the change of unemployment rate could be more helpful in making the prediction than the change of gold price. Inspired by the Tensorflow Quick, Draw! tutorial [7] [1], we believe societal events can be represented as composite sketches, where the characteristics for the sketch are the statistics of various indicators of the events. Given these indicators, we can reconstruct the sketch in hope of identifying future events.

The 2 categories of data are modeled by the following approaches - sequences in terms of "characteristic strokes" for structured data, and weighted average vectors for unstructured text data, as illustrated by Figure 2.

4.3.1 Structured data. Indicator values are represented as $1 \times d$ vectors, where d is the size of input window. Values are then normalized, and differences between consecutive days are calculated.

4.3.2 Unstructured text data. For tweets, we use the publicly available word vectors trained on Google News[8] [10], which contains 300-dimensional vectors for 3 million words and phrases. We refine the pre-trained word vectors by using a separate set of tweets collected within the civil unrest domain (tweets which contain the keywords in the dictionary). We also translate Arabic tweets into English. Therefore, tweets are cleaned by removing URLs, preserved words like "RT", etc., and represented as 30×300 matrices, where the first 30 words in tweets are considered based on the median length of cleaned tweets in our dataset is 15. Weighted average matrices of all the tweets from the input window of size d are then fed as input into the framework, where weights are calculated as user impact [7] and normalized to 0-to-1 range.

$$S_{user,date} = \ln\left(\frac{(follower + \theta)^2}{following + \theta}\right), \tag{1}$$

where θ is a smoothing constant set equal to 1 so that natural logarithm is always applied on a real positive number. As the number of followers and followings may change over times, the user impact score is calculated for a given user on a given date. It is expected that users with higher number of followers and fewer followings are more popular and influential in the real world.

[7]Source: https://www.tensorflow.org/tutorials/sequences/recurrent_quickdraw
[8]https://code.google.com/archive/p/word2vec/

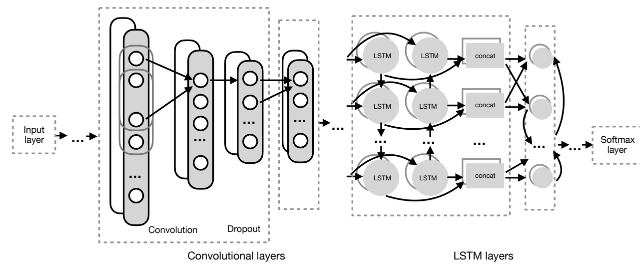

Figure 3: The framework of our model.

4.4 Predictive Models

Clear shortcomings of traditional feedforward neural networks have been proven when being applied to time series related tasks, because time sequences are hard to be captured by these models. LSTM Networks are a special kind of recurrent neural network which addresses this issue and has the capability to learn long term dependencies in the input sequences. Convolutional neural networks are proposed to be used as encoder as it has been convincingly shown that they can produce a rich representation of the input [15]. Based on the belief that civil unrest events unfold complex mechanisms and are the comprehensive results of multiple factors, we propose an encoder-decoder like model, Cov-LSTM, which uses a combination of convolutional layers and LSTM layers, to learn patterns from various data sources and predict the CU events effectively.

Figure 3 shows the architecture of our model. Input layer takes in vector representations of structured data or unstructured text data. A series of 1-dimensional convolutional layers if structured input, or 2-dimensional layers if unstructured text input are added as well as the dropout layers. Output of the convolutional layer is then input into a stack of bidirectional LSTM layers. Finally a softmax layer is used as the output layer for the classification task.

5 EXPERIMENTAL RESULTS

5.1 Experimental Setup

We conduct experiments on 2 datasets - GSR data and ACLED data, and focus on country-level predictions for Egypt and Jordan, city-level predictions for Cairo, Amman, and Delhi. Models are trained separately for each country/city. Each data point corresponds to a day and is labeled by the number of events happened on that day. Days with more than some certain number of events are categorized as one class based on the observation that most of the data points fall into the first classes. Daily values which are updated on weekly/monthly basis are calculated by assigning the same value to the days in that week/month. Sparse softmax cross entropy is used to measure the probability error when training and testing the models. To find hyper-parameters, we use 10-fold cross validation.

5.2 Evaluation Metrics

Performance is judged by taking into account of both the lead time (at least 3 days) and the Quality Score (QS). Lead time is defined as the average number of days between the date the prediction is generated and the reported event date. Quality Score measures the difference between the predicted counts and actual counts, and is calculated as Equation 2. Average is taken among all the data points being tested.

$$Quality\ Score = 1 - \frac{abs(Predicted - Actual)}{max(Predicted, Actual, 4)} \tag{2}$$

The Mercury Challenge uses the Mercury Score for ranking, which is calculated as follows:

$$Mercury\ Score = 1,000,000 * QS \tag{3}$$

Autoregressive Integrated Moving Average (ARIMA) model is used as the benchmark algorithm. ARIMA model combines the properties of autoregressive models, where future values of the series are determined by the tendency of the series to revert to the mean following one or more shocks, and moving-average models, where future values of the series are predicted to be the weighted average of preceding values.

5.3 Results

5.3.1 Country-level prediction. Table 2 summarizes the Mercury Score of our model in comparing with the benchmark algorithm provided by the challenge, and two other leading groups on the leaderboard, "rekcahd" and "valilenk". Unfortunately, we currently can not obtain more information about the algorithms used by the competitors.

Days with more than 10 events are classified into one class, i.e., 12 classes in total, as around 90% of the days in GSR dataset have fewer than 11 events. Input features are selected by recursive feature elimination. Social media indicators and economic indicators are then used. Feature values from the previous 30 days are used to forecast the event counts for 3 days later. We use 3 layers of 1-dimensional convolutional layer with (48, 64, 96) filters of length (5, 5, 3), and 3 layers of bidirectional LSTM layer with 128 nodes per layer. A dropout rate of 0.3 is added to the output of each layer.

The results shown in Table 2 demonstrate that by using social media and selected economic indicators, our model outperforms the other competing groups as well as the benchmark algorithm, ARIMA, trained based on the pre-challenge GSR historical data. This also proves that models will need to be retrained with new data in order to learn the most up-to-date information for tasks involving time series data. Specifically, for predictions made for Egypt, our proposed model beats the base rate model in every evaluation period. Though it is hard to beat the base rate, our model shows strong advantages over the other competing groups for all evaluation periods.

Table 2: Mercury Scores calculated for each evaluation period of our model (Cov-LSTM), base rate model (ARIMA), and top-ranking groups.

	2018-08-01 to 2018-10-31	2018-11-01 to 2018-11-13	2018-11-01 to 2018-11-27	2018-11-01 to 2018-12-11	2018-11-01 to 2018-12-25	2018-11-01 to 2019-01-08
Egypt						
ARIMA	316479	269231	388889	390244	372727	362319
Cov-LSTM	798007	**846153**	787037	**829268**	**831818**	**829710**
rekcahd	**835015**	836152	**789545**	772238	802162	806661
valilenk	823034	788462	731481	762195	768182	778986
Jordan						
ARIMA	634114	**1000000**	**888889**	**758608**	**730986**	**699858**
Cov-LSTM	**658532**	647435	698765	680081	691060	699348
rekcahd	521188	768555	679934	571972	615693	647649
valilenk	378845	500000	515873	595238	612670	572817

5.3.2 City-level prediction. We conduct evaluation on the ACLED dataset over 2 months (January, February, 2019) for 3 cities: Cairo, Amman, and Delhi. We include a third city Delhi as it is the city where a large number of events are recorded by ACLED, while the other two cities, there are no events in most of the days. We implement Rolling Forecast ARIMA Model using historical data to show the advantage of our approach.

We use 3 classes to categorize data on Cairo and Amman, 5 classes for Delhi. As around 99% of the days have fewer than 2 events for Cairo and Amman. For Delhi, more than 90% days have fewer than 4 events. To use Twitter embeddings as input, 3 2d-convolutional layers with (16, 32, 48) filters of size (5, 5, 3) are used, followed by a 3-layer bidirectional LSTM with 64 nodes per layer. Dropout rate is set to be 0.3. Performance of historical ACLED event data based model and Twitter data embedding model under multiple settings are compared with ARIMA (5, 1, 0). AUC-ROC graphs are shown in Figure 4 and Figure 5.

The results demonstrate that our framework outperforms the benchmark algorithm with configurations (choice of input window size and lead time). In general, better performance can be achieved when using data from the previous 30 days (blue lines) than 60 days (red lines). Therefore, including more historical data is not necessarily helpful for improving the prediction accuracy. We can also observe that Twitter data embedding model performs well even with increased lead time, which indicates that including dynamic data, i.e. tweets, can potentially be able to help increase the prediction lead time. The model performance for Cairo and Amman are better than for Delhi as more predicting classes are used for Delhi.

6 CONCLUSIONS

Clues to Civil Unrest events can be discerned in advance through indicators of economical, political, or social conditions. Such indicators often change rapidly and vary from country to country.

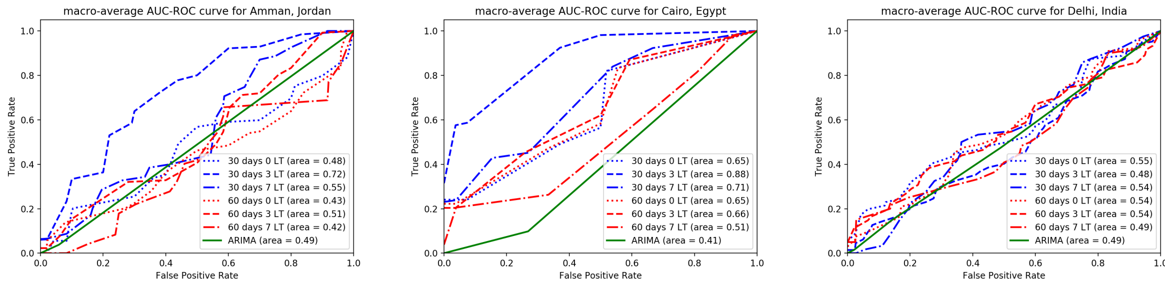

(a) Amman, Jordan **(b) Cairo, Egypt** **(c) Delhi, India**

Figure 4: ROC curves of historical (ACLED event) data based model with different lead time (LT) for varying input window sizes.

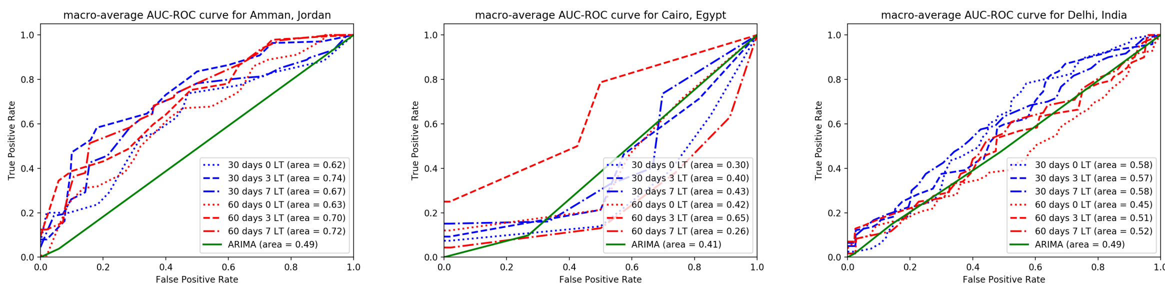

(a) Amman, Jordan **(b) Cairo, Egypt** **(c) Delhi, India**

Figure 5: ROC curves of Twitter data embedding model on ACLED data with different lead time (LT) for varying input window sizes.

We see promising results by including heterogenous data sources for predicting civil unrests and using LSTM models that effectively exploit sequential data. Our approach demonstrates remarkable results on both city-level and country-level predictions with structured and unstructured data. This work shows that it is possible to leverage existing datasets to provide predictions of civil unrest with sufficient lead time and granularity to be used in early warning systems for tasks such as effective resource allocation, which could help reduce the economic losses and the impact on people's daily lives.

REFERENCES

[1] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2015. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. https://www.tensorflow.org/ Software available from tensorflow.org.

[2] Jose Cadena, Gizem Korkmaz, Chris J Kuhlman, Achla Marathe, Naren Ramakrishnan, and Anil Vullikanti. 2015. Forecasting social unrest using activity cascades. *PloS one* 10, 6 (2015), e0128879.

[3] Taylor Dewey, Juliane Kaden, Miriam Marks, Shun Matsushima, and Beijing Zhu. 2012. The impact of social media on social unrest in the Arab Spring. *International Policy Program* 5 (2012).

[4] Andrew Hoegh, Scotland Leman, Parang Saraf, and Naren Ramakrishnan. 2015. Bayesian model fusion for forecasting civil unrest. *Technometrics* 57, 3 (2015), 332–340.

[5] Nathan Kallus. 2014. Predicting crowd behavior with big public data. In *Proceedings of the 23rd International Conference on World Wide Web*. ACM, 625–630.

[6] Gizem Korkmaz, Jose Cadena, Chris J Kuhlman, Achla Marathe, Anil Vullikanti, and Naren Ramakrishnan. 2015. Combining heterogeneous data sources for civil unrest forecasting. In *Proceedings of the 2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2015*. ACM, 258–265.

[7] Vasileios Lampos, Nikolaos Aletras, Daniel Preoţiuc-Pietro, and Trevor Cohn. 2014. Predicting and characterising user impact on Twitter. In *14th conference of the European chapter of the Association for Computational Linguistics 2014, EACL 2014*. 405–413.

[8] Kalev Leetaru and Philip A Schrodt. 2013. Gdelt: Global data on events, location, and tone, 1979–2012. In *ISA annual convention*, Vol. 2. Citeseer, 1–49.

[9] Pedro Manrique, Hong Qi, Ana Morgenstern, Nicolas Velasquez, Tsai-Ching Lu, and Neil Johnson. 2013. Context matters: improving the uses of big data for forecasting civil unrest: emerging phenomena and big data. In *Intelligence and Security Informatics (ISI), 2013 IEEE International Conference on*. IEEE, 169–172.

[10] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*. 3111–3119.

[11] Fengcai Qiao and Kedi Chen. 2016. Predicting Protest Events with Hidden Markov Models. In *Cyber-Enabled Distributed Computing and Knowledge Discovery (CyberC), 2016 International Conference on*. IEEE, 109–114.

[12] Fengcai Qiao, Pei Li, Xin Zhang, Zhaoyun Ding, Jiajun Cheng, and Hui Wang. 2017. Predicting social unrest events with hidden Markov models using GDELT.

Discrete Dynamics in Nature and Society 2017 (2017).

[13] Clionadh Raleigh, Andrew Linke, Håvard Hegre, and Joakim Karlsen. 2010. Introducing acled: An armed conflict location and event dataset: Special data feature. *Journal of peace research* 47, 5 (2010), 651–660.

[14] Naren Ramakrishnan, Patrick Butler, Sathappan Muthiah, Nathan Self, Rupinder Khandpur, Parang Saraf, Wei Wang, Jose Cadena, Anil Vullikanti, Gizem Korkmaz, et al. 2014. 'Beating the news' with EMBERS: forecasting civil unrest using open source indicators. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 1799–1808.

[15] Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. 2015. Show and tell: A neural image caption generator. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 3156–3164.

[16] Frank DW Witmer, Andrew M Linke, John O?Loughlin, Andrew Gettelman, and Arlene Laing. 2017. Subnational violent conflict forecasts for sub-Saharan Africa, 2015–65, using climate-sensitive models. *Journal of Peace Research* 54, 2 (2017), 175–192.

[17] Congyu Wu and Matthew S Gerber. 2017. Forecasting Civil Unrest Using Social Media and Protest Participation Theory. *IEEE Transactions on Computational Social Systems* (2017).

[18] Mirae Yang. 2012. The Collision of Social Media and Social Unrest: Why Shutting Down Social Media is the Wrong Response. *Nw. J. Tech. & Intell. Prop.* 11 (2012), xix.

[19] Liang Zhao, Qian Sun, Jieping Ye, Feng Chen, Chang-Tien Lu, and Naren Ramakrishnan. 2015. Multi-task learning for spatio-temporal event forecasting. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 1503–1512.

[20] Liang Zhao, Jieping Ye, Feng Chen, Chang-Tien Lu, and Naren Ramakrishnan. 2016. Hierarchical incomplete multi-source feature learning for spatiotemporal event forecasting. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 2085–2094.