

EMBERS at 4 years: Experiences operating an Open Source Indicators Forecasting System

Sathappan Muthiah¹, Patrick Butler¹, Rupinder Paul Khandpur¹,
Parang Saraf¹, Nathan Self¹, Alla Rozovskaya¹, Liang Zhao¹, Jose Cadena²,
Chang-Tien Lu¹, Anil Vullikanti², Achla Marathe², Kristen Summers^{*}, Graham Katz[‡],
Andy Doyle[‡], Jaime Arredondo[¶], Dipak K. Gupta^{||}, David Mares[¶], Naren Ramakrishnan¹

¹ Discovery Analytics Center, Dept. of Comp. Science, Virginia Tech, Arlington, VA 22203

² Biocomplexity Institute, Virginia Tech, Blacksburg, VA 24061

[¶] University of California at San Diego, San Diego, CA 92093

^{||} San Diego State University, San Diego, CA 92182

[‡] CACI Inc., Lanham, MD 20706

^{**} IBM Watson Group, Chantilly, VA 20151

ABSTRACT

EMBERS is an anticipatory intelligence system forecasting population-level events in multiple countries of Latin America. A deployed system from 2012, EMBERS has been generating alerts 24x7 by ingesting a broad range of data sources including news, blogs, tweets, machine coded events, currency rates, and food prices. In this paper, we describe our experiences operating EMBERS continuously for nearly 4 years, with specific attention to the discoveries it has enabled, correct as well as missed forecasts, lessons learnt from participating in a forecasting tournament, and our perspectives on the limits of forecasting including ethical considerations.

Keywords

civil unrest, event forecasting, open source indicators.

1. INTRODUCTION

Modern communication forms such as social media and microblogs are not only rapidly advancing our understanding of the world but also improving the methods by which we can comprehend, and even forecast, the progression of events. Tracking population-level activities via ‘massive passive’ data has been shown to quite accurately shed light into large-scale societal movements.

Two years ago, in KDD 2014, we described EMBERS [15], a deployed anticipatory intelligence system [5] that forecasts significant societal events (e.g., civil unrest events such as protests, strikes, and ‘occupy’ events) using a large set of

open source indicators such as news, blogs, tweets, food prices, currency rates, and other public data. The EMBERS system has been running continuously 24x7 for nearly 4 years at this point and our goal in this paper is to present the discoveries it has enabled, both correct as well as missed forecasts, and lessons learned from participating in a forecasting tournament including our perspectives on the limits of forecasting and ethical considerations. In particular, we shed insight into the value proposition to an analyst and how EMBERS forecasts are communicated to its end-users.

The development of EMBERS is supported by the Intelligence Advanced Research Projects Activity (IARPA) Open Source Indicators (OSI) program. EMBERS currently focuses on multiple regions of the world but for the purpose of this paper we focus primarily on the 10 Latin American countries, specifically the countries of Argentina, Brazil, Chile, Colombia, Ecuador, El Salvador, Mexico, Paraguay, Uruguay, and Venezuela. Similarly, EMBERS generates forecasts for multiple event classes—fluenza-like-illnesses [3], rare diseases [16], elections [12], domestic political crises [8], and civil unrest—but in this paper we focus primarily on civil unrest as this was the most challenging event class with hundreds of events every month across the countries studied here. EMBERS forecasts are scored against the Gold Standard Report (GSR), a monthly catalog of events as reported in newspapers of record in these 10 countries. The GSR is compiled by MITRE corporation using human analysts.

Our key contributions can be summarized as follows:

1. Unlike retrospective studies of predictability, EMBERS forecasts are communicated in real-time before the event to MITRE/IARPA and scored independently of the authors. (Further, the scoring criteria are set by IARPA.) We present multiple quantitative indicators of EMBERS performance as well as insights into how we made EMBERS forecasts most valuable to analysts. We report two primary ways in which analysts utilize EMBERS and present the use of automated narratives to help make EMBERS forecasts as useful as possible.

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 ACM 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.

KDD '16, August 13-17, 2016, San Francisco, CA, USA

© 2016 ACM. ISBN 978-1-4503-4232-2/16/08...\$15.00

DOI: <http://dx.doi.org/10.1145/2939672.2939709>

- In an attempt to demystify the state-of-the-art in forecasting and to create an open dialogue in the community, we report both successful forecasts of EMBERS as well as events missed by EMBERS. The events not forecast by EMBERS lead us to considerations of both the limitations of the underlying technology as well as the inherent limits to forecasting large-scale events.
- While social media is often touted as the key to event forecasting systems such as EMBERS, we present the results of an ablation study to outline the performance degradation that ensues if data sources such as Twitter and Facebook were to be removed from the forecasting pipeline.
- We consider the separation of civil unrest events into events that happen with a degree of regularity versus rare or significant events, and evaluate the performance of EMBERS in forecasting such surprising events.
- We describe our current best understanding of the limitations to forecasting civil unrest events using technologies like EMBERS and also consider the ethical considerations of the EMBERS technologies.

2. BACKGROUND

We begin by providing a brief review of forecasting systems, followed by a quick preview of EMBERS, its system architecture, machine learning models, and measures for evaluating its performance. For more details, please see [15].

Forecasting societal events such as civil unrest has a long tradition in the intelligence analysis and political science community. We distinguish between forecasting systems versus event coding systems (systems that provide structured representations of ongoing events reported in newspapers), and focus on the former. Early forecasting systems such as ICEWS [14] provided very broad coverage in countries but were limited by their spatio-temporal resolution (e.g., typically country- and month- level forecasting for specific events of interest [19]). The ICEWS events of interest are domestic political crises, international crises, ethnic/religious violence, insurgencies, and rebellion. A similar project in scope is PITF (Political Instability Task Force) [6] funded by the CIA. To the best of our knowledge, only EMBERS provides the most-specific spatial resolution (city-level) and the most-specific temporal resolution (daily-level) capability in forecasting.

The software architecture of EMBERS (Early Model Based Event Recognition using Surrogates) is designed as a loosely coupled, share-nothing, highly distributed pipeline of processes connected via ZeroMQ. In this manner, the system is both highly scalable and fault tolerant. The EMBERS pipeline can loosely be broken up into four stages: ingestion, enrichment, modeling, and selection. In the first stage, ingestion, data is collected from a variety of sources and streamed into the following stages in real-time. The enrichment stage takes the raw data from the ingestion stage and processes it in various ways including natural language processing, geocoding, and relative time phrase normalization. After enrichment, the modeling stage feeds the enriched data into the various models that make up EMBERS. Unlike other systems which use single monolithic models to make predictions, EMBERS combines the results of several different models to arrive at the most accurate forecasts.




Figure 1: An example depicting how an alert is scored with respect to the ground truth.




Figure 2: Alert sent at time t_1 predicting an event at time t_3 can be matched to a GSR event that happened at time t_2 and reported at time t_4 if $t_1 < t_4$.

In particular, the separate alerts from each model are deduplicated, fused, and selected and finally emitted as a full forecast for a real world event.

The structure of a civil unrest forecast is shown in Figure 1 (left). A forecast constitutes four fields, corresponding to the when, where, who, and why of the protest. These fields are respectively denoted as the date, location, population, and event type. Location is recorded at the city level. Population and event type are fields chosen from a categorical set of possibilities. The figure further shows how an alert with all these fields are scored against a GSR event. In the basic scoring methodology shown in Figure 1 each of the four fields are weighted uniformly and a total quality score out of 4 is obtained. Apart from this each alert also has a lead-time associated with it calculated as shown in Figure 2.

Rather than design one model to integrate all possible data sources, EMBERS adopted a multi-model approach to forecasting. Each model utilized a specific (possibly overlapping) set of data sources and is tuned for high precision, so that the union of these models can be tuned for high recall. A fusion/suppression engine [7] allows a tunable strategy to issue more or fewer alerts depending on whether the analyst’s objective is to obtain a higher precision or recall. The underlying models used in EMBERS are: (i) *planned protest model* [13], (ii) *dynamic query expansion* [20], (iii) *volume-based model* [9], (iv) *cascade regression* [2], and (v) a baseline model. The planned protest model, for news and social media (Twitter, Facebook), identifies explicit signs of orga-

nization and calls for protest, resolves relative mentions of time (e.g., ‘next Saturday’) and space (e.g., ‘the square’) to issue forecasts. The dynamic query expansion (DQE) model uses Twitter as a data source and learns time- and country-specific expansions of a seed set of keywords to identify specific situational circumstances for civil unrest. For instance, in Venezuela (an economy where the government exercises stringent price controls), there were a series of protests in 2014 stemming from the shortage of toilet paper, a novel circumstance that was uncovered by DQE. The volume-based model uses a range of data sources, spanning social, economic and political indicators. It uses classical statistical models (LASSO and hybrid regression models) to forecast civil unrest events using features from social media (Twitter and blogs), news sources, political event databases (ICEWS and GDELT [10]), Tor [4] statistics, food prices, and currency exchange rates. It aims to provide a multi-source perspective into forecasting by leveraging the selective superiorities of different data sources. The cascade regression model aims to model activity related to organization and mobilization in Twitter [2]. Finally, the baseline model uses maximum likelihood estimation over the GSR to issue history-based forecasts.

The EMBERS project is unique not just in its algorithmic underpinnings but also in the use of new measures for evaluation, specifically aimed at determining forecasting performance. As shown in Figure 2, one of the primary measures of EMBERS performance is lead time, the number of days by which a forecast ‘beats the news’, i.e., the date of reporting of the event. Lead time should not be confused with date quality, i.e., the difference between the predicted date and the actual date of the event. The date quality is one of the components of the quality score, the other components being the location score, event type score, and population score. Figure 1 shows how these other components are scored between an EMBERS forecast and a GSR record.

Given a set of alerts and a set of GSR events for a given month, the lead time is used as a constraint to define legal (alert, event) pairs so that we can construct a bipartite matching to optimize the best quality score. From this bipartite matching, measures of precision and recall can be derived, i.e., by assessing the number of (un)matched events or alerts. Finally, a confidence score is used to assess the quality of probabilities imputed by EMBERS to its forecasts, and measured in terms of the Brier score. For more details, please see [15].

We now turn to a discussion of specific discoveries enabled by EMBERS, into civil unrest in Latin America, and into the complexity of the forecasting enterprise as a whole.

3. PERFORMANCE ANALYSIS

First, we begin with a performance analysis of EMBERS, from both a quantitative point of view with respect to the GSR and with respect to end-user (analyst) goals.

3.1 Quantitative Metrics

Figure 3 depicts both the targets set by the IARPA OSI program as well as the actual measures achieved by the EMBERS system. As shown here, the easiest target to achieve in EMBERS was, surprisingly, the lead time objective. This was feasible due to EMBERS’s focus on modeling both planned and spontaneous events. Planned events are sometimes organized with as many as several weeks of

Targets			
	Month 12	Month 24	Month 36
Mean Lead-Time	1 day	3 days	7 days
Mean Probability Score	0.60	0.70	0.85
Mean Quality Score	3.0	3.25	3.5
Recall	0.50	0.65	0.80
Precision	0.50	0.65	0.80

Actual			
Metric	Month 12	Month 24	Month 36
Mean Lead-Time	3.89 days	7.54 days	9.76 days
Mean Probability Score	0.72	0.89	0.88
Mean Quality Score	2.57	3.1	3.4
Recall	0.80	0.65	0.79
Precision	0.59	0.94	0.87

Figure 3: IARPA OSI targets and results achieved by EMBERS.




Figure 4: Comparison of number of perfect scores (4.0) obtained by EMBERS vs a baserate model each month in 2013.

lead time and thus identifying indicators of organization was instrumental in achieving lead time objectives. The confidence (mean probability) scores were also achieved by EMBERS and involved careful calibration of probabilities by taking into account estimates of model propensities and data source reliabilities. The measure that was most difficult to achieve was the quality score as it involved a four component additive score and thus tangible improvements in score required more than incremental improvements in forecasting specific components. Finally, recall and precision involve a natural underlying trade-off and the deployment of our fusion/suppression engine provided the ability to balance this trade-off to meet IARPA OSI’s objectives. Apart from comparing mean scores another interesting measure is to see how many perfect matches (4.0 quality score) are obtained by EMBERS. Fig. 4 shows the number of alerts issued by EMBERS that matched perfectly to an event in the future on a monthly basis for 2013. It is clear that EMBERS makes almost double the number of fully accurate forecasts as compared to a baserate model. The baserate model generates alerts using the rate of occurrence of events in the past three months.

3.2 Analyst Evaluation

In addition to the quantitative measures above, our experience interacting with analysts demonstrated an interesting dichotomy as to how analysts use EMBERS alerts. Some analysts preferred to use EMBERS in an ‘analytic triage’ scenario wherein they could tune EMBERS for high recall so that they would apply their traditional measures of filtering

EMBERS forecasts that there will be a **violent** protest on **February, 18th 2014** in **Caracas**, the **capital city of Venezuela**. It predicts that the protest will involve people working in the **business sector**. The protest will be related to **discontent about economic policies**.

There were **5, 5, and 5** other similar warnings in last **2, 7 and 30** days, respectively.

The forecast date of the warning falls in **week 7**, which may have **historical importance**; this **week is found to be statistically significant** ($pval=0.00461919415894$, $zscore=2.832$, $avg. count=57.25$, $mean=21.569 \pm 12.597$)

Audit trail of the warning includes an **article printed 2014-02-17**.

Major players involved in the protest include **Venezuelan opposition leader, students, President Nicolas Maduro, and Leopoldo Lopez**.

Reasons: Protest **against rising inflation and crime**; Protestors want a **political change**; President Nicolas Maduro has **accused US consular officials** and **right-wing**.

Protests are characterized by: **Venezuelan opposition leader spearheaded days of protest** and **calling for peaceful demonstration**; **Maduro accused official on 2014-12-16**; Protests have seen **several deadly street protests**; **Three people were killed on 2014-02-12**; **Demonstrations setting days of clashes**; **supporters to march to Interior Ministry on 2014-02-18**.

Figure 5: An example narrative for an EMBERS alert. Here, color **red** indicates named entities, **green** refers to descriptive protest related keywords. Items in **blue** are historical or real time statistics and those in **magenta** refer to inferred reasons of the protest.

and analysis to hone in on forecasts of interest. Other analysts instead viewed EMBERS as a data source and preferred to use it in a high precision mode, e.g., wherein they were focused on a specific region of the world (e.g., Venezuela) and/or aimed to investigate a particular social science hypothesis (e.g., whether disruptions in global oil markets led to civil unrest).

To support these diverse classes of users, we implemented two mechanisms in the alert delivery stage. First, we implemented a mechanism wherein in addition to generating alerts, EMBERS also forecasted the expected quality score for each forecast (using machine learning methods trained on past GSR-alert matches). This expected quality score measure provided a way for analysts to use quality directly as a way to tune the system to receive greater or fewer alerts. Second, we implemented an automated narrative generation capability (see Fig. 5) wherein EMBERS auto-generates a summary of the alert in English prose. As shown in Fig. 5, a narrative comprises many parts drawn from different sources of information. One source constitutes the named entities wherein the system uses ‘Wikification’ to identify definitions and descriptions of named entities on Wikipedia. A second source is historical (or real-time) statistics of warning output and warning performance and situating the alert in this context. The third source pertains to inferred reasons for the protest using knowledge graph identification techniques.

4. EMBERS SUCCESSES AND MISSES

Next, we detail some of the successful as well as not so successful forecasts made by EMBERS over the past few years in Latin America.




Figure 6: EMBERS performance during the Brazilian Spring (June 2013).

4.1 Successful Forecasts

Brazilian Spring (June 2013)

These protests were the largest and most significant protests in Brazil’s recent history and caught worldwide attention. Millions of Brazilians took part in these demonstrations, also known as the Brazilian Spring or the Vinegar Movement (inspired from the use of vinegar soaked cloth by demonstrators to protect themselves from police teargas). These protests were sparked by an increase in public transport fares from R\$3 to R\$3.20 by the government of President Dilma Rousseff.

As shown in Fig. 6, while missing the initial uptick, EMBERS did forecast the increase in the order-of-magnitude of protest events during the Brazilian Spring and also captured the spatial spread in the events. In addition EMBERS correctly forecast that this event will span the broad Brazilian general population (as opposed to being confined to specific sectors).

Around 68% of EMBERS alerts during this period originated from the planned protest model. This is due to the fact that social networking platforms (Twitter and Facebook) as well as conventional news media played a key role in organization of these uprisings. Although initial protests were primarily due to the bus fare increases, they quickly morphed into more broader dissatisfaction to include wider issues such as government corruption, over-spending, and police brutality. The demonstrators also made calls for political reforms. In response, President Rousseff proposed a plebiscite on widespread political reforms in Brazil (but this was later abandoned). Through its dynamic query expansion model, EMBERS was able to capture such discussions on Twitter (see Fig. 7), and tracked their evolution as events unfolded through June.

The protests intensified in late June (see Fig. 6), which were forecast correctly by EMBERS, and these events also coincided with FIFA 2013 Confederation Cup matches. We believe this was an important factor in helping the protests gain momentum, as the events were covered by international media. A majority of protests occurred in the cities that were hosting FIFA soccer matches. EMBERS issued most of its alerts for these host cities (see Fig. 8), viz. Rio de Janeiro, São Paulo, Belo Horizonte, Salvador, and Porto Alegre, among others. For example, on 27th June during the Confederations Cup semi-final in Fortaleza, around 5000




Figure 11: Timeline of Mexico protests, showing the correspondence between counts of GSR events and EMBERS alerts on a daily basis.

Colombia protests (Dec 2014 to March 2015)

Colombia witnessed two different significant protests during this period, one during late December 2014 and the other during February 2015. Towards the end of 2014, the Colombian government was on the process of moving forward with peace negotiations to end 50 years of conflict with the Revolutionary Armed Forces of Colombia (FARC). With the FARC rebels having been associated with various acts of terror, e.g., extortion, armed conflict, kidnapping, ransom, and illegal mining, for a long period, the people of Colombia gathered in huge numbers to protest against possible amnesty for the FARC rebels. EMBERS successfully forecast the uptick in the number of events during the middle of December 2014 as indicated in Fig. 12. This figure also depicts the increase in protest counts during February 2015, although in this case EMBERS over-predicted the counts.

The protests in February 2015 were qualitatively different in nature and were led by truckers unions demanding better freight rates, labor rights, and revolting against high fuel prices. The truckers protests extended for about a month and caused an estimated loss of about \$300 million to the Colombian economy. EMBERS forecast the truckers protests accurately at the onset of these events but over-estimated the number of protests during February 11-12.

Paraguay protests (February 2015)

The February 2015 protests in Paraguay were mainly carried out by peasants against the actions of President Horacio Cartes. The protests were carried out after president Horacio's public revelation that he had opened two private Swiss bank accounts. The protests also had a historical significance. They were also being carried out as a tribute to peasant leaders and activists who were murdered. The peasants also protested against the introduction of a new public-private partnership law. EMBERS forecast the uptick in the number of Paraguay protests during mid February 2015 as shown in Fig. 13.




Figure 12: EMBERS performance during the Colombia protests (Dec 2014 to Mar 2015).




Figure 13: EMBERS performance during the Paraguay protests (Feb. 2015).




Figure 14: EMBERS performance during the Brazilian protests of March 2015.

4.2 EMBERS Misses

Next, we outline specific large-scale events that EMBERS failed to forecast accurately, along with a discussion of underlying reasons.

4.2.1 Brazilian protests (March 2015)

The beginning of 2015 saw a series of protests in Brazil demanding the removal of president Dilma Roussef amidst much furor against the increasing corruption in the country. The number of protests increased significantly due to the revelations that many politicians belonging to the ruling party accepted bribes from the state-run energy company Petrobras. The protests drew huge participation from the general population, with protesters generally estimated to be around a million. EMBERS, as shown in Figure 14, picked up the onset of events but failed to capture the sudden rise in the number of events.

During this period there was a significant architectural change in the EMBERS processing pipeline. As mentioned in Section 2 the EMBERS system enrichment pipeline consists of the following steps: natural language processing, geocoding, and relative time phrase normalization (temporal tagging). During early 2015 EMBERS had moved to the Heideltime [17] temporal tagger versus the previously used TIMEN [11] temporal tagger. (This choice was made because Heideltime supported more languages and an active development cycle.) Heideltime had no support for Portuguese (the primary language of Brazil) and the EMBERS software development team had extended Heideltime to support Portuguese by translating the underlying resources for Spanish to Portuguese. As it turns out, the simple translation of rules from Spanish to Portuguese was not sufficient and this affected the recall of one of the key models for Brazil, viz. the planned protest model. Since the planned protest model relies almost exclusively on the quality of information (specifically, date) extraction from text, its performance significantly deteriorated. This was subsequently corrected for the future by adding more rules and correcting existing rules (which were translated from Spanish) in Heideltime for Portuguese with the aid of language experts and extensive backtesting. Fig. 15 shows an example detection before and after the changes.

"Os manifestantes voltará novamente em catorze dias"	
Publication Date: 2015-02-01	
BEFORE	AFTER
<pre>{"datetimes": []}</pre>	<pre>{"datetimes": [{"expr": "catorze dias", "normalized": "2015-02-15"}]}</pre>

Figure 15: Example depicting the improvement in time phrase recognition after changes to Heideltime.




Figure 16: EMBERS performance during Mexico protests (Oct 2014).

4.2.2 Mexico protests (Dec 2014)

The days of December 2014 witnessed a continuation of the series of protests that began in October 2014 as described in Section 4.1. People turned out in huge numbers in different cities of Mexico demanding President Pena Nieto's ouster owing to the manner in which the case of the 43 missing students were handled. The protests were largely peaceful except for a few cases where vehicles were torched and windows and office equipment were broken.

EMBERS missed the huge single day spike on December 1st. Though having predicted a nationwide event for December 1, EMBERS failed to capture the individual cities where the protests would take place out and thus was unable to forecast the number of events accurately. On manual retrospective, it was found that the date, viz. December 1, was picked by the protesters due to its historical significance. This was the day when President Pena Nieto was sworn in (in 2012) amidst much controversy and opposition from many specific constituent groups. The manual analysis also led to the understanding of how special dates were mentioned by the twitterati #1Dmx. Dates mentioned using such abbreviations went unrecognized by the EMBERS system and was one of the main reasons for EMBERS not being able to capture the peak on December 1st despite its historical significance.

4.2.3 Brazilian Spring Onset

The EMBERS forecasts during the 2013 Brazilian Spring as shown in Fig. 6 was able to capture the peak but as can be seen the system was unable to capture the initial onset. See Section 7 for a detailed discussion of the limits of forecasting.




Figure 17: EMBERS ablation visualizer.

Table 1: Comparison of performance measures under ablation testing. Social media sources contribute toward recall but, due to their noisy nature, lower other measures of performance.

Data Source	Quality-Score	Lead-time	Precision	Recall
Removing news and blogs	-16.48%	-55%	+35%	-14%
Removing social media	+8.42%	+30%	+79%	-33%

5. ABLATION TESTING

Different data sources provide different value to the forecasting enterprise. It is important that we understand the value of a data source w.r.t. its forecasting potential. In this section we describe ablation testing in EMBERS where the incremental value addition is evaluated for specific data sources. In particular, we are interested in determining the utility of using social media versus traditional media such as news and blogs. Table 1 shows the percentage improvement or degradation in performance measures when specific sources are removed. It clearly shows that social media sources are mainly necessary in achieving high recall but are not that useful in achieving high lead times, for which traditional media sources are required. This behavior is expected as social media is where daily chatter occurs whereas signs of organization and calls for protest often happen over (or are reported in) traditional media. Mainly, Table 1 makes it evident that to build a successful forecasting system we need a good mix of both traditional and social media sources. Fig. 17 shows a snapshot of the EMBERS ablation visualizer. The visualizer provides an analyst with the capability to selectively remove data sources and assess the differences in final alerts.

6. FORECASTING SURPRISING EVENTS

The GSR contains a mix of everyday, mundane, protests as well as surprising events such as the Brazilian Spring. We aimed to ascertain the relative ease of forecasting each class of events with respect to the baserate model described in section 3.1. To define surprising events, we employed a maximum entropy approach. For this purpose, each event is assumed to be describable in terms of three dimensions: country, population, and event type. The GSR can then be conceptualized as a cube. We infer a maximum entropy

Table 2: Surprising events identified by a maximum entropy analysis over the GSR.

Month	Country	Description
May '13	Colombia	Nationwide protests against administrative policy changes regarding the Youth in Action program and social security pension.
Jun '13	Brazil	Brazilian Spring.
Jul '13	Brazil	Brazilian Spring.
Sep '13	Mexico	Nationwide protests against education and energy reforms.
Oct '13	Mexico	Nationwide protests against education and energy reforms.
Oct '13	Uruguay	Nationwide protests demanding increase in minimum wage.
Feb '14	Venezuela	Venezuelan student protests.
Mar '14	Venezuela	Venezuelan student protests.
May '14	Brazil	Nationwide demonstrations in response to the 2014 FIFA World Cup and other social issues.
Jun '14	Brazil	Nationwide demonstrations in response to the 2014 FIFA World Cup and other social issues.
Aug '14	Argentina	Nationwide protests against drops in wages, employment, and rising inflation.
Sep '14	Ecuador	Nationwide protests to demand changes in labor policies.
Oct '14	Mexico	Nationwide protests after the discovery of mass graves of kidnapped students.

distribution conditioned on the marginals induced by the cube using iterative proportional fitting [1].

Every month, events from the last three months of the GSR are used to populate the underlying cube of counts and an iterative proportional fitting procedure is used to estimate the expected counts for each cell. The resultant sample counts are then scaled to match the observed number of GSR events in the current month. The cell-wise difference between the inferred maxent distribution and the observed GSR is computed and all cells with significance greater than five standard deviations are classified as containing surprising events. In essence, this approach takes the GSR as input and creates a truncated GSR against which we can evaluate EMBERS (and the baserate model).

In Table 2 we present several inferred events in Latin America from our maximum entropy filter. For all these events, we compare the recall of both EMBERS and baserate model in Figure 18. It is clear that EMBERS is able to forecast these significant upticks consistently and that the baserate model is not able to.

7. UNCERTAINTIES IN FORECASTING

While examining the events of civil unrest closely in the past few years, it was clear to our team that events carry two distinct types of uncertainties: **cause** and **timing**. Fig. 19 summarizes these uncertainties.

Among all the incidents of civil unrest that we encounter, the largest and the most significant ones are planned events. These events are usually organized by political parties, labor and student unions. Since it takes a huge effort to organize protest demonstrations that attract thousands, the organizers must disseminate information regarding the venue and




Figure 18: Performance of EMBERS vs a baserate model for surprising events.

		Cause	
		Known	Unknown
Timing	Known	Planned Events	Recurring Events
	Unknown	Spontaneous Events	Black Swan Events

Figure 19: Studying the limitations of event forecasting.

the date and time. These announcements are posted on the organizers' websites and are widely shared on social media. By scouring our sources, it has been possible for EMBERS to accurately forecast the occurrence of these types of protests.

The recurring events take place on a regular basis. For instance, in Chile and Argentina the 'mothers of the disappeared' protest the disappearance of their children by the military dictatorships of the 1970s and 1980s on a certain particular day of the week and in the same plaza. In some countries with large Muslim populations, fighting and protests break out regularly after Friday evening prayer as people stream out of the mosques after listening to fiery sermons. These are typically small events but if they are reported as part of our GSR, EMBERS models will be able to forecast them.

The protests for which the causes are known but not the timing are staged spontaneously. These events are the outcomes of longstanding frustration and anger which fuel widespread protests in response to trigger events. Thus, the viral videos of police brutality or a sudden change in government policy can start a prairie fire of protests. The Brazilian Spring with origins in bus fares and which channeled public anger against corruption and government mismanagement is a classical example. The challenge here is not just to be aware of the underlying tensions that might erupt when an event occurs, but to also distinguish between events that do and do not perform as triggers. Algorithms to better model precursors is an area of further research that will aid further forecasting this class of events.

Finally, *black swan* events [18] are rare and truly unforeseen and can happen as a result of natural disasters, the

sudden death of a leader, or even the sudden rise of a small group that can truly destabilize a nation. (Each of these types of events can in turn spark civil unrest.) For instance the rise of the Islamic State in Iraq and Syria (ISIS) has truly confounded policymakers all over the world. While there were other Sunni groups, from al-Qaeda to al-Nusra, that contributed to instability, the rapid ascendance of ISIS, which did not depend on an isolated terrorist attack and burst out with a clear holding of territories as a full-scale insurgency, surprised most observers. It might not be feasible to forecast the beginnings of such events; however, once such movements have been initiated, models should be able to detect and forecast their momentum.

8. ETHICAL ISSUES

EMBERS, as an anticipatory intelligence system, has many powerful legitimate uses but is also susceptible to abuse.

First, it is important to have a discussion of civil unrest and its role in society. In our opinion, under the proper circumstances, civil unrest enhances the ability of citizens to communicate not only their views but also their priorities to those who govern them. Governments constantly need to make choices and find it difficult to know, on specific issues at particular times, how their constituencies value the available options. Elections are retrospective indicators and rarely issue-specific; polling taps into sentiment, but is not a good indicator of priorities or strength of feeling because of the low cost associated with responding. Events, on the other hand, indicate a willingness to bear some costs (organization, mobilization, identification) in support of an issue and thus reveal not only preferences but provide some indication of priorities.

An open sources indicators approach, as we have used here, is a potentially powerful tool for understanding the social construction of meaning and its translation into behavior. EMBERS can contribute to making the transmission of citizen preferences to government less costly to the economy and society as well. There are economic costs to even peaceful disruptions embodied in civil unrest due to lost work hours and the deployment of police to manage traffic and the interactions between protestors and bystanders. Given the vulnerability of large gatherings to provocation by handfuls of violence-oriented protestors (e.g., Black Box anarchists in Brazil) the economic, social and political costs of large-scale public demonstrations are also potentially significant to marchers, bystanders, property owners and the government – democratically elected or not. The right to demonstrate can still be respected but if the government responds to grievances in time, the protestors may cancel the event or fewer people might participate in the event. In today's interconnected society, protests also cause disruptions to supply chain logistics, travel, and other sectors, and anticipating disruptions is key to ensuring safety as well as reliability.

The potential power of civil unrest forecasting systems, like those of most scientific advances, is susceptible to abuse by both democratic and non-democratic governments. The appropriate safeguards require developing transparent and accountable democratic systems, not outlawing science. Non-democratic governments may clearly abuse such forecasting systems. But even here the value of forecasting civil unrest is not simply negative. Many non-democratic regimes transition to democratic ones, often in a violent process but not

always (in Latin America, authoritarian regimes negotiated transitions to democracy without a civil war in Mexico, Honduras, Peru, Bolivia, Brazil, Uruguay, Argentina and Chile). The rational choice models of authoritarian decision-making in such crises always explain a dictatorship’s collapse rather than accommodation to a transition by pointing to the lack of credible information in a dictatorship regarding citizens’ true feelings. EMBERS-like models may thus provide the information that facilitates and encourages some transitions.

9. CONCLUSIONS

We have presented our experiences and lessons learnt from operating EMBERS continuously 24x7 over a period of four years. EMBERS has proven itself to be a reliable predictor of civil unrest events in 10 different countries by ingesting data from multiple languages. EMBERS has been a wonderful testbed for our data science team and even in its misses much has been learnt.

Future work falls primarily in three directions. First, we intend to investigate more formal methods to model precursors to protests, so that better precursor detection can constitute better protest forecasting. Second, we will continue to build a strong systems capability to forecasting societal events, by enlarging the scope of regions and languages covered. Finally, we are investigating techniques to remove or reduce the human element required in generating the GSR. Currently the most human intensive part of the EMBERS project is generating a GSR for training and validation of the models.

Acknowledgments

Supported by the Intelligence Advanced Research Projects Activity (IARPA) via DoI/NBC contract number D12PC000337, the US Government is authorized to reproduce and distribute reprints of this work for Governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DoI/NBC, or the US Government.

10. REFERENCES

- [1] Y. M. Bishop, S. E. Fienberg, and P. W. Holland. *Discrete multivariate analysis: theory and practice*. Springer Science & Business Media, 2007.
- [2] J. Cadena, G. Korkmaz, C. J. Kuhlman, et al. Forecasting Social Unrest using Activity Cascades. *PLoS One*, 10(6):e0128879, 2015.
- [3] P. Chakraborty, P. Khadivi, B. Lewis, A. Mahendiran, et al. Forecasting a Moving Target: Ensemble Models for ILI Case Count Predictions. In *SIAM International Conference on Data Mining, April 24-26, 2014*, pages 262–270, 2014.
- [4] R. Dingledine, N. Mathewson, and P. F. Syverson. Tor: The Second-Generation Onion Router. In *USENIX Security Symposium, August 9-13, 2004*, pages 303–320, 2004.
- [5] A. Doyle, G. Katz, K. Summers, C. Ackermann, et al. Forecasting Significant Societal Events using The EMBERS Streaming Predictive Analytics System. *Big Data*, 2(4):185–195, dec 2014.
- [6] J. A. Goldstone, R. H. Bates, D. L. Epstein, et al. A Global Model for Forecasting Political Instability. *American Journal of Political Science*, 54(1):190–208, 2010.
- [7] A. Hoegh, S. Leman, P. Saraf, and N. Ramakrishnan. Bayesian Model Fusion for Forecasting Civil Unrest. *Technometrics*, 57(3):332–340, February 2015.
- [8] Y. Keneshloo, J. Cadena, G. Korkmaz, and N. Ramakrishnan. Detecting and Forecasting Domestic Political Crises: A Graph-Based Approach. In *ACM Web Science Conference, WebSci ’14, June 23-26, 2014*, pages 192–196, 2014.
- [9] G. Korkmaz, J. Cadena, C. J. Kuhlman, A. Marathe, A. Vullikanti, and N. Ramakrishnan. Combining Heterogeneous Data Sources for Civil Unrest Forecasting. In *IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, ASONAM, August 25 - 28, 2015*, pages 258–265, 2015.
- [10] K. Leetaru and P. Schrodt. GDELT: Global Data on Events, Location and Tone, 1979–2012. *ISA Annual Convention*, pages 1979–2012, 2013.
- [11] H. Llorens, L. Derczynski, et al. TIMEN: An Open Temporal Expression Normalisation Resource. In *International Conference on Language Resources and Evaluation, LREC, May 23-25, 2012*, pages 3044–3051, 2012.
- [12] A. Mahendiran, W. Wang, J. A. S. Lira, et al. Discovering Evolving Political Vocabulary in Social Media. In *International Conference on Behavioral, Economic, and Socio-Cultural Computing, BESEC, October 30 - November 1, 2014*, pages 26–32, 2014.
- [13] S. Muthiah, B. Huang, J. Arredondo, et al. Planned Protest Modeling in News and Social Media. In *AAAI Conference on Artificial Intelligence, January 25-30, 2015*, pages 3920–3927, 2015.
- [14] S. P. O’Brien. Crisis Early Warning and Decision Support: Contemporary Approaches and Thoughts on Future Research. *International Studies Review*, 12(1):87–104, 2010.
- [15] N. Ramakrishnan, P. Butler, S. Muthiah, N. Self, et al. ‘Beating the news’ with EMBERS: Forecasting Civil Unrest using Open Source Indicators. In *International Conference on Knowledge Discovery and Data Mining, KDD, August 24 - 27, 2014*, pages 1799–1808, 2014.
- [16] T. Rekatsinas, S. Ghosh, S. R. Mekaru, et al. SourceSeer: Forecasting Rare Disease Outbreaks using Multiple Data Sources. In *SIAM International Conference on Data Mining, April 30 - May 2, 2015*, pages 379–387, 2015.
- [17] J. Strötgen and M. Gertz. HeideTime: High Quality Rule-based Extraction and Normalization of Temporal Expressions. In *International Workshop on Semantic Evaluation, SemEval ’10*, pages 321–324, 2010.
- [18] N. N. Taleb. *The Black Swan. The Impact of the Highly Improbable*. Random House Inc., 2008.
- [19] M. D. Ward, N. W. Metternich, C. Carrington, et al. *Geographical Models of Crises: Evidence from ICEWS*, pages 429–438. CRC Press, Boca Raton, FL, 2012.
- [20] L. Zhao, F. Chen, et al. Unsupervised Spatial Event Detection in Targeted Domains with Applications to Civil Unrest Modeling. *PLoS ONE*, 9(10):e110206, 2014.