

Predicting Crowd Behavior with Big Public Data

Nathan Kallus
Massachusetts Institute of Technology
77 Massachusetts Ave E40-149
Cambridge, MA 02139
kallus@mit.edu

ABSTRACT

With public information becoming widely accessible and shared on today's web, greater insights are possible into crowd actions by citizens and non-state actors such as large protests and cyber activism. We present efforts to predict the occurrence, specific timeframe, and location of such actions before they occur based on public data collected from over 300,000 open content web sources in 7 languages, from all over the world, ranging from mainstream news to government publications to blogs and social media. Using natural language processing, event information is extracted from content such as type of event, what entities are involved and in what role, sentiment and tone, and the occurrence time range of the event discussed. Statements made on Twitter about a future date from the time of posting prove particularly indicative. We consider in particular the case of the 2013 Egyptian coup d'état. The study validates and quantifies the common intuition that data on social media (beyond mainstream news sources) are able to predict major events.

Categories and Subject Descriptors

I.2.1 [Computing Methodologies]: Artificial Intelligence—*Applications and Expert Systems*

Keywords

Web and social media mining, Twitter analysis, Crowd behavior, Forecasting, Event extraction, Temporal analytics, Sentiment analysis, Online activism

1. INTRODUCTION

The manifestation of crowd actions such as mass demonstrations often involves collective reinforcement of shared ideas. In today's online age, much of this public consciousness and coming together has a significant presence online where issues of concern are discussed and calls to arms are publicized. The Arab Spring is oft cited as an example of the new importance of online media in the formation of mass

protests [6]. While the issue of whether mobilization occurs online is highly controversial, that nearly all crowd behavior in the internet-connected world has some presence online is not. So while it may be infeasible to predict an action developing in secret in a single person's mind, the ready accessibility to public information on the web that future crowds may now be reading and reacting to or members of which are now posting on social media can offer glimpses into the formation of this crowd and the action it may take.

News from mainstream sources from all over the world can now be accessed online and about 500 million tweets are posted on Twitter each day with this rate growing steadily [15]. Blogs and online forums have become a common medium for public discourse and many government publications are offered for free online. We here investigate the potential of this publicly available information online for predicting mass actions that are so significant that they garner wide mainstream attention from around the world. Because these are events perpetrated by human actions, they are in a way endogenous to the system, enabling prediction.

But while all this information is in theory public and accessible and could lead to important insights, gathering it all and making sense of it is a formidable task. We here use data collected by Recorded Future (www.recordedfuture.com). Scanning over 300,000 different open content web sources in 7 different languages and from all over the world, mentions of events—in the past, current, or said to occur in future—are continually extracted at a rate of approximately 50 extractions per second. Using natural language processing, the type of event, the entities involved and how, and the timeframe for the event's occurrence are resolved and made available for analysis. With such a large dataset of what is being said online ready to be processed by a computer program, the possibilities are infinite. For one, as shown herein, the gathering of crowds into a single action can often be seen through trends appearing in this data far in advance.

We here study the cases of mass protests and politically motivated cyber campaigns involving an entity of interest, such as a country, city, or organization. We use historical data of event mentions online, in particular forward-looking mentions of events yet to take place, to forecast the occurrence of these in the future. We can make predictions about particular timeframes in the future with high accuracy.

We find that the mass of publicly available information online has the power to unveil the future actions of crowds. Measuring the trends in sheer numbers we are here able to accurately predict protests in countries and in cities and large cyber campaigns by target and by perpetrator. After

assembling our prediction mechanism for significant protests, we investigate the case of the 2013 Egyptian coup d'état and how well we were able to foresee the protests surrounding it.

The ability to forecast these things has important ramifications. Countries and cities faced with a high likelihood of significant protest can prepare themselves and their citizens to avoid any unnecessary violence and damages. Companies with personnel and supply chain operations can ask their employees to stay at home and to remain apolitical and can attempt to safeguard their facilities in advance. Countries, companies, and organizations faced with possible cyber campaigns against them can beef up their cyber security in anticipation of attacks or even preemptively address the cause for the anger directed at them.

Recent work has used online media to glean insight into consumer behavior. In [2] the authors mine Twitter for insights into consumer demand with an application to forecasting movie earnings. In [8] the authors use blog chatter captured by IBM's WebFountain [7] to predict Amazon book sales. These works are similar to this one in that they employ very large data sets and observe trends in crowd behavior by huge volumes. Online web searches have been used to describe consumer behavior, most notably in [3] and [5], and to predict movements in the stock market in [4].

In [14] the authors study correlations between singular events with occurrence defined by coverage in the New York Times. By studying when does one target event ensue another specified event sometime in the future, the authors discover truly novel correlations between events such as between natural disasters and disease outbreaks. Here we are interested in the power of much larger, more social, and more varied datasets in pointing out early trends in endogenous processes (actions by people discussed by people) that can help predict all occurrences of an event and pinning down *when* they will happen, measuring performance with respect to each time window for prediction. One example of the importance of a varied dataset that includes both social media and news in Arabic is provided in the next section.

We here seek to study the predictive power of such web intelligence data and not simply the power of standard machine learning algorithms (e.g. random forests vs. SVMs). Therefore we present only the learning machine (random forest in the case of predicting protests) that performed best on the training data and compare it to a data-poor maximum-likelihood random-walk predictor that predicts for the future the situation today. The data used in this study has been made available for download at www.nathankallus.com/PredictingCrowdBehavior/.

We first discuss what predictive signals we expect to find in order to motivate our constructions. We then review how mentions of events, entities, and times are extracted from the wide breadth of sources. Using this data we develop a predictive mechanism to predict significant protests in countries and in cities. We consider the case of the 2013 Egyptian coup d'état and conclude. In an appendix [9] we consider finding more general patterns in the data, motivating an application of the naïve Bayes classifier to high-dimensional sequence mining in massive datasets, which we use to forecast cyber attack campaigns by target or perpetrator.

2. PREDICTIVE SIGNALS IN PUBLIC DATA

We begin by exemplifying anecdotally the precursory signals that exist in public data for large protests. On Sunday

June 9, 2013 a Beirut protest against Hezbollah's interference in Syria turned violent when clashes with Hezbollah supporters left one protester dead [12]. The story was widely reported on June 9 including in Western media, attracting more mainstream news attention than any protest event in Lebanon in over a year marking it as a *significant protest*.

But not only were there signs that the protest would occur before it did, there were signs it may be large and it may turn violent. The day before, Algerian news source Ennahar published an article with the headline "Lebanese faction organizes two demonstrations tomorrow rejecting the participation of Hezbollah in the fighting in Syria" (translated from Arabic using Google Translate). There was little other preliminary mainstream coverage and no coverage (to our knowledge) appeared in mainstream media outside of the Middle-East-North-Africa (MENA) region or in any language other than Arabic. Moreover, without further context there would be little evidence to believe that this protest, if it occurs at all, would become large enough or violent enough to garner mainstream attention from around the world.

However, already by June 5, four days earlier, there were many Twitter messages calling people to protest on Sunday, saying "Say no to #WarCrimes and demonstrate against #Hezbollah fighting in #Qusayr on June 9 at 12 PM in Downtown #Beirut" and "Protest against Hezbollah being in #Qusair next Sunday in Beirut." In addition, discussion around protests in Lebanon has included particularly violent words in days prior. A June 6 article in TheBlaze.com reported, "Fatwa Calls For Suicide Attacks Against Hezbollah," and a June 4 article in the pan-Arabian news portal Al Bawaba reported that, "Since the revolt in Syria, the security situation in Lebanon has deteriorated." A May 23 article in The Huffington Post mentioned that, "The revolt in Syria has exacerbated tensions in Lebanon, which ... remains deeply divided." Within this wider context, understood through the lens of web-accessible public information such as mainstream reporting from around the world and social media, there was a significant likelihood that the protest would be large and turn violent. These patterns persist across time; see Figure 1.

It is these predictive signals that we would like to mine from the web and employ in a learning machine in order to forecast significant protests and other crowd behavior. In this case, it is critical that we spread a wide net on many sources so to catch mentions in non-Western media and foreign-language Tweets along with mentions in media (such as Reuters) and languages (such as English) with a more global reach.

3. EVENT EXTRACTION

To quantify these signals we will look at time-stamped event-entity data. The data harvesting process extracts




Figure 1: Number of same-day news reports of protest in Lebanon (red) and tweets mentioning protests said to occur over the next three days (blue). Red spikes often follow blue spikes or otherwise have a recognizable convex ramp-up.

mentions of events from the plethora of documents continually gathered from the over 300,000 sources being monitored. An important aspect is that the event mentions are tagged with the time range in which the event is said to occur in the mention so that forward-looking statements, such as plans to protest, can be directly tied to a future time and place. Event extractions are done in Arabic, English, Farsi, French, Russian, Spanish, and Simplified and Traditional Chinese.

There are several elements in the event, entity, and time extraction process. For each document mined from the web, an ensemble of off-the-shelf natural language processing tools are used to extract tokens (lemma, root, stem, and part of speech) and entities. Entities extracted by each tool are then combined and resolved into a database of canonical entities for disambiguation. Documents are categorized by topic. Entities that are not agreed upon by the various tools and are far-fetched given the topic are rejected. Ontologies of structured entity relationships constructed from online sources (such as DBpedia) are used to guide filtering and provide a gazetteer for additional entity extraction.

Given the set of filtered entities, a statement extractor links these entities to events stated in the document. Events are again extracted using an ensemble method of off-the-shelf tools and a custom made tool that relies on the above tokenization. Each event from the various extraction sources is matched up with a particular text fragment from the document that best represents it. Next, n -grams in the document are matched against phrase lists organized by sentiment or tone and the fraction of these in the fragment is recorded.

All time statements made in the document are separately extracted. The tokenization of the document is parsed by a dependency grammar using the data-driven parser Malt-Parser [11] to construct a dependency graph of the document. This is used to find time statements, both relative (e.g. “next summer”) and absolute. Comprehending these in machine time is based on several contextual cues. Cultural and regional categorizations are extracted from the document to inform such things as relevant hemisphere for seasons, which is the first day of the week, standard date formats (month first or day first), and timezone. Moreover, contexts such as publishing date are noted. Using these, all time statements made in the document are converted to standardized time-stamps, with specificity varying from second to year. Event mentions are then matched up with the most relevant time statement to the event statement based on sentence dependency.

In all, events are marked by type of event, time range of event, entities involved, role of entities involved, entities mentioned, sentiment and tone, and origin. Some post-processing is done on the event-entity level to further improve quality based on special curated ontologies. For example, known hacker groups such as Anonymous constitute one such ontology and if mentioned in a cyber event but not clearly as the perpetrator these are assumed to be so nonetheless. Similarly, impossible events are rejected. For example, one ontology keeps track of death dates of people based primarily on information harvested from Wikipedia and assists in rejecting the event “Marco Polo will travel to China in 2015” as impossible because Marco Polo is dead.

The precision of event extraction is measured by workers on the Amazon Mechanical Turk (mturkers). To test the precision of time-stamping, for each language language-specific Human Intelligence Tasks (HITs) are constructed for

Country	All	Twitter	Country	All	Twitter
Afghanistan	60918	27655	Lebanon	44153	23394
Bahrain	246136	177310	Libya	162721	69437
Egypt	944998	397105	Nigeria	70635	38700
Greece	122416	70521	Pakistan	289643	213636
India	491475	274027	S. Arabia	39556	13670
Indonesia	34007	17120	Sudan	28680	13654
Iran	118704	53962	Syria	212815	79577
Italy	65569	43803	Tunisia	99000	27233
Jordan	35396	19369	Yemen	70583	16712

Table 1: Protest event mentions in the corpus.

a random sample. Each HIT consists of the text fragment, the extracted time-stamp, and the question “is extracted time right or wrong?” and is given to three mturkers. It is declared successful if at least two answer “right.” For example, the precision of time-stamping in both English and Spanish is measured at 93%, in Arabic at 90%, and in Simplified Chinese at 82%. The precision of the event extraction is measured similarly by type. Protest events in English come in at 84%. Malware threat events in English come in at 96% and in Simplified Chinese at 90%.

4. PREDICTING SIGNIFICANT PROTESTS

We now turn our attention to the use of this event data to the prediction of significant protests around the world. Our first forecasting question will revolve around predicting significant protests on the country level and considering that country alone. That is, a significant protest is one that receives much more same-day mainstream reporting than is usual for that country. So while most days a business with supply chain operations in Egypt operate under the usual volatile circumstances (since 2011) of Egypt—certainly more volatile than, say, Jordan and receiving much more attention for it—they are interested in receiving advance notice of protests that are going to be larger and more dangerous than the ordinary for Egypt. The same for another milieu. At the same time, we will use past patterns from other countries to inform the prediction mechanism when making a prediction about one country. In fact, the prediction mechanism will not be knowledgeable of the particular country in question but instead just the *type* of country, quantified by cluster membership. In an appendix [9] we also investigate the question of predicting protests that are significant relative to a common baseline and on the city level instead.

We restrict to a selection of 18 countries: Afghanistan, Bahrain, Egypt, Greece, India, Indonesia, Iran, Italy, Jordan, Lebanon, Libya, Nigeria, Pakistan, Saudi Arabia, Sudan, Syria, Tunisia, and Yemen. We will consider all protest event mentions in any of these countries being published any time between January 1, 2011 and July 10, 2013 as the event mention corpus. January 1, 2011 up to March 5, 2013 will serve for supervised training (and validation by cross-validation) and since March 6, 2013 up to July 10, 2013 will serve for test and performance scoring. Let

$$M_{cs}(i, j) = \begin{array}{l} \text{Number of event mentions of protest in} \\ \text{country } c \text{ taking place on day } j \text{ extracted} \\ \text{from documents published on day } i \text{ from} \\ \text{sources of type } s \end{array}$$

We tabulate some totals of these numbers over the whole event mention corpus in Table 1. For example, nearly one million mentions of a protest event in Egypt occur in the data, over a third of a million on Twitter.




Figure 2: Normalized count of mainstream reports $M''(i)$ in Egypt with annotations for stretches above $\theta = 2.875$ (red line).

4.1 The Ground Set

For each country, the protests we are interested in forecasting are those that are significant enough to garner more-than-usual real-time coverage in mainstream reporting for the country. That is, there is a significant protest in country c on day i if $M_{c,\text{Mainstream}}(i, i)$ is higher than usual for country c (in an appendix [9] we consider an absolute scale for protest significance which naturally boosts predictability). Since new sources are being added daily by the hundreds or thousands to the Recorded Future source bank, there is a heterogeneous upward trend in the event mention data and what is more than usual in counts changes. To remove this trend we normalize the mention counts by the average volume in the trailing three months. That is, we let

$$M'_{cs}(i, i+k) = \frac{M_{cs}(i, i+k)}{\frac{1}{|\text{Countries}| \times 90} \sum_{c' \in \text{Countries}} \sum_{j=i-90}^{i-1} M_{c's}(j, j+k)}$$

where Countries is the 18 countries. Next we define the training-set average of same-day mainstream reporting

$$\bar{M}'_c = \frac{1}{|\text{Train}|} \sum_{i \in \text{Train}} M'_{c,\text{Mainstream}}(i, i)$$

where Train denotes the set of days in the training set.

Moreover, to smooth the data we consider a three-day moving average. Then we say, by definition, that a significant protest in country c (and relative to country c) occurs during the days $i-1, i, i+1$ if

$$M''_c(i) = \frac{1}{3} \sum_{j=i-1}^{i+1} \frac{M'_{c,\text{Mainstream}}(j, j)}{\bar{M}'_c} \geq \theta$$

is larger than a threshold θ . The threshold is chosen so to select only significant protests. By inspecting the data's correspondence to the largest protests, we set $\theta = 2.875$ (which is also nearly the 94th percentile of the standard exponential distribution). Overall across all countries considered, this resulted in 6% of 3-day stretches to be labeled positive, distributed mostly evenly among the countries.

An example plot of $M''_c(i)$ for Egypt is shown in Figure 2 with annotations of top mainstream news describing the protests in each stretch of above-threshold days. Notable protest events provide the reader with anchor points are the 9/11 anniversary protests in 2012 (concurrent with the

Benghazi attacks), the late-June protests leading up to the Egyptian coup d'état, the riots set off by soccer-fan violence in early 2012, and the riots after 30 of the fans involved were sentenced to death which also coincided with riots connected to the anniversary of the revolution in early 2013.

4.2 Scoring Protest Predictions

We are interested in predicting on each day i whether a significant protest will occur over the next three days $i+1, i+2, i+3$ based on information published on or before i . That is, on each day i we wish to predict whether $M''_c(i+2) \geq \theta$ (which depends on days $i+1, i+2, i+3$).

We quantify the success of a predictive mechanism based on its balanced accuracy. Let $T_{ci} \in \{0, 1\}$, $P_{ci} \in \{0, 1\}$ respectively denote whether a significant protest occurs in country c during the days $i+1, i+2, i+3$ and whether we predict there to be one. The true positive rate (TPR) is the fraction of positive instances ($T_{ci} = 1$) correctly predicted to be positive ($P_{ci} = 1$) and the true negative rate (TNR) is the fraction of negative instances predicted negative. The balanced accuracy (BAC) is the unweighted average of these: $\text{BAC} = (\text{TPR} + \text{TNR})/2$. BAC, unlike the marginal accuracy, cannot be artificially inflated. Always predicting "no protest" without using any data will yield a high 94% marginal accuracy but only 50% balanced accuracy. In fact, a prediction without any relevant data will always yield a BAC of 50% on average by statistical independence.

4.3 The Features

We now attempt to quantify the predictive signals we encountered anecdotally in Section 2. These features will serve as the data based on which we make predictions.

In Section 2 we exemplified how the violence in language surrounding discussion of protest in a country can help set the context for the danger of a future protest to get out of hand. Each event mention is rated for violent language by the fraction of n -grams in the corresponding fragment (ignoring common words) that match a phrase list. Let

$$V_{cs}(i) = \text{Total violence rating of fragments associated with event mentions of protest in country } c \text{ extracted from documents published on day } i \text{ from sources of type } s$$

Similarly to the normalization of event mentions due to the

ever-growing source bank, we normalize this quantity as

$$V'_{cs}(i) = \frac{V_{cs}(i)}{\frac{1}{|\text{Countries}| \times 90} \sum_{c' \in \text{Countries}} \sum_{j=i-90}^{i-1} V_{c's}(j)}$$

In addition, forward-looking mentions in mainstream news and Twitter can help indicate whether a protest is planned and estimate how many might plan to attend. We have already defined $M_{cs}(i, i+k)$ which counts this data for $k \geq 1$.

In order to facilitate trans-country training, we normalize these features with respect to the series we would like to predict, $M'_{c, \text{Mainstream}}(i, i)$. Similar to the normalization of $M''_c(i)$, we normalize these features by a per-country constant coefficient $(\overline{M'_c})^{-1}$.

For the purposes of trans-country training, we hierarchically cluster the countries using Ward's method [17] applied with the distance between two countries c, c' equal to the Kolmogorov-Smirnov uniform distance between the distribution functions of the set of training values of M'' ignoring the time dimension. That is, $d(c, c') =$

$$\sup_{z \geq 0} \left| \frac{1}{|\text{Train}|} \sum_{i \in \text{Train}} (\mathbb{I}\{z \geq M''(i)\} - \mathbb{I}\{z \geq M''(i')\}) \right|$$

This distance is also a non-parametric test statistic to test the hypothesis that two samples were drawn from the same distribution. We construct $\lfloor 2\sqrt{|\text{Countries}|} \rfloor$ clusters using R function `hclust` [13]. We include as a feature the indicator unit vector of cluster membership of the country c associated with the instance (c, i) . Thus the classifier does not know the particular country about which it is making a prediction, just its type as characterized by this clustering.

For each instance (c, i) we also include as features the ten most recent days of same-day reporting on protest in c ,

$$\frac{M'_{c, \text{Mainstream}}(i, i)}{\overline{M'_c}}, \dots, \frac{M'_{c, \text{Mainstream}}(i-9, i-9)}{\overline{M'_c}}$$

along with the two most recent differences of these values. We also include the violence rating in recent mainstream reporting as the cumulative partial sums of the values

$$\frac{V'_{c, \text{Mainstream}}(i)}{\overline{M'_c}}, \dots, \frac{V'_{c, \text{Mainstream}}(i-9)}{\overline{M'_c}}$$

Next we include the counts of mentions of protests said occur over the next three days, published either in mainstream news or in Twitter over the ten recent days. We incorporate this feature as the cumulative partial sums of the values

$$\frac{\sum_{k=1}^3 M'_{cs}(i, i+k)}{\overline{M'_c}}, \dots, \frac{\sum_{k=1}^3 M'_{cs}(i-9, i+k)}{\overline{M'_c}}$$

for $s = \text{Mainstream}$ and $s = \text{Twitter}$.

When predicting farther into the future, about the three days starting with the k^{th} day from today, we push all indexes back by $k-1$ days, thus excluding any future data and maintaining the same overall length of the feature vector.




Figure 3: BAC for predicting protests by distance into the future using the random forest and the full data-set (blue and trend in dashed black) and the data-poor predict-like-today heuristic (red).




Figure 4: Achievable TPR-FPR rates as we vary the classification threshold for (left) country-level predictions with a relative significance scale and (right) city-level predictions with an absolute significance scale.

4.4 The Classifier

For the prediction mechanism we employ a random forest classifier trained on all data up to March 5, 2013. We use the R library `randomForest` [10]. The parameters are left at their default values as set in the library. For example, the defaults dictate that the forest have 500 trees each trained on $\lfloor \sqrt{\# \text{features}} \rfloor$ randomly chosen features.

We tune only the threshold required for a positive prediction. If the fraction of trees in the forest voting positive is at least this threshold then a positive prediction is made, otherwise negative. We tune this by four-fold cross validation over the training set using the BAC metric.

4.5 Results

We tested the trained random forest on test data March 6, 2013 to July 10, 2013. The results were $\text{TPR}=\mathbf{75.51\%}$ and $\text{TNR}=\mathbf{69.31\%}$. BAC, the average of these two, is $\mathbf{72.41\%}$, which constitutes a $\mathbf{44.8\%}$ reduction in balanced error from having no data. In comparison, predicting for the future the situation today, simulating the data-poor prediction possible when one nonetheless has information about today's situation (whether from being at the location or from news), has $\text{TPR}=\mathbf{27.04\%}$, $\text{TNR}=\mathbf{93.74\%}$, and $\text{BAC}=\mathbf{60.39\%}$.

As we attempt to predict farther into the future our predictions become noisier and closer to the no data case as the earlier data has less bearing on the far future and there are fewer reports mentioning events to occur on the days in question. In Figure 3 we plot the accuracy of making predictions farther into the future. For each $k \geq 1$ we re-train the random forest with the pushed-back feature vectors.

As we vary the voting threshold, we can (monotonically) trade off true positives with true negatives. We plot the range of achievable such rates for our classifier in Figure 4 (left), using the false positive rate ($\text{FPR} = 1 - \text{TNR}$) as is common by convention. The area under the curve (AUC), here $\mathbf{78.0\%}$, is the probability that a random positive instance has more positive votes than a random negative instance. The fraction of trees voting positive on an instance could well serve as a risk rating. By randomizing between two thresholds, we can achieve any weighted average of the rates. The rates achievable via randomization is the convex hull of the curve in Figure 4 and has $\mathbf{79.2\%}$ under it.

To show the possible breadth of this approach, we can also consider an alternative forecasting question where we instead wish to predict protests on the city level and we define significance on a global scale (see appendix [9] for more details). That is, a protest is deemed significant if it garners unusually high same-day mainstream reporting relative to




Figure 5: Predictions of protests in Egypt around the coup d'état. Yellow to red mark positive predictions and blue to purple negative. Redder colors indicate more positive votes.

all cities considered. Using a global scale, the fraction of trees voting positive serves as an absolute risk rating that is comparable between cities. The predictability of positive events also increases because knowing the city alone is no longer statistically independent of there being a positive event (for example, 23% of positive training instances are in Cairo). The further localized data also improves predictability. On the test data, the random forest yields $\text{TPR}=84.7\%$, $\text{TNR}=85.7\%$, $\text{BAC}=85.2\%$. The achievable rates as the voting threshold is varied are shown in Figure 4 (right) with $\text{AUC}=91.3\%$ and 91.9% for the convex hull.

4.6 The Case of the Egyptian Coup d'État

To exemplify the prediction mechanism we study predictions made for Egypt on different days about the relative future around the time of the 2013 coup d'état in Egypt (Figure 5). As shown, the days around June 30 were predicted positive with very high certainty for a long time prior, the date ranges in 6/28–7/2 being consistently predicted positive since June 6, three weeks beforehand (since June 16 onward shown in figure). Indeed, with a lot of discontent with and talk of demonstration against President Morsi's rule, many protests were anticipated for the weekend of June 30, the anniversary of Morsi's rise to the presidency. Even U.S. Secretary of State John Kerry made a statement in anticipation of protests asking for peaceful and responsible demonstration on, as he says, Saturday (June 29) and Sunday (June 30) [16]. Therefore, those days were long predicted positive with high certainty. However, already on June 28 spontaneous "warm-up" protests burst in the streets [1]. The first range to include June 28 was predicted with less certainty, especially from farther back, but was correctly predicted starting June 10 except on June 16 when it was mistakenly reported negative, just slightly below the threshold.

As we now know the protests around the anniversary did grow very large with many injured and dead in clashes with police and demonstrators from opposite camps. The protests continued and on July 1 the Egyptian army issued an ultimatum to Morsi to resolve the protests within 48 hours or it would intervene. On July 3 the Egyptian army removed Morsi from power. Protests intensified and many more people died. As seen in the figure, already on June 28 when the protests had only just started, before the anniversary and before any discussion of a possible ultimatum or coup, the prediction mechanism had already correctly declared that significant protests will go on for the weeks to come.

5. CONCLUSIONS

We presented new findings about the power of massive online-accessible public data to predict crowd behavior. With much of public discourse having at least some presence online and usually more, the wide range of public data captured by our efforts offers unparalleled insight into the futures of countries, cities, and organizations as affected by mass demonstrations and cyber campaigns (see appendices [9]). The evidence presented validates and quantifies the common intuition that data on social media (beyond mainstream news sources) are able to predict major events. The scope and breadth of the data offered glimpses into tweets in foreign languages and news in far places. The confluence of all this information showed trends to come far in the future. We are excited to make the data used in this study public and invite further exploration into its predictive abilities.

6. REFERENCES

- [1] Ahram.org. [Egypt warms up for a decisive day of anti- and pro-Morsi protests](#). Accessed: 2013-08-25.
- [2] S. Asur and B. A. Huberman. Predicting the future with social media. In *WI-IAT*, 2010.
- [3] H. Choi and H. Varian. Predicting the present with google trends. *Economic Record*, 88(s1):2–9, 2012.
- [4] Z. Da, J. Engelberg, and P. Gao. In search of attention. *J. Finance*, 66(5):1461–1499, 2011.
- [5] S. Goel, J. M. Hofman, S. Lahaie, D. M. Penneck, and D. J. Watts. Predicting consumer behavior with web search. *PNAS*, 107(41):17486–17490, 2010.
- [6] S. González-Bailón, J. Borge-Holthoefer, A. Rivero, and Y. Moreno. The dynamics of protest recruitment through an online network. *Sci. Rep.*, 1, 2011.
- [7] D. Gruhl, L. Chavet, D. Gibson, J. Meyer, P. Pattanayak, A. Tomkins, and J. Zien. How to build a WebFountain: An architecture for very large-scale text analytics. *IBM Syst. J.*, 43(1):64–77, 2004.
- [8] D. Gruhl, R. Guha, R. Kumar, J. Novak, and A. Tomkins. The predictive power of online chatter. In *SIGKDD*, 2005.
- [9] N. Kallus. Predicting Crowd Behavior with Big Public Data (preprint with appendix). 2014. <http://arxiv.org/abs/1402.2308>.
- [10] W. Liaw. Classification and regression by randomForest. *R News*, 2(3):18–22, 2002.
- [11] J. Nivre, J. Hall, J. Nilsson, A. Chanev, G. Eryigit, S. Kübler, S. Marinov, and E. Marsi. MaltParser: A language-independent system for data-driven dependency parsing. *Nat. Lang. Eng.*, 13(2):95–135, 2007.
- [12] NYTimes.com. [Protester Dies in Clash That Apparently Involved Hezbollah Supporters](#). Accessed: 2013-08-24.
- [13] R Core Team. *R: A Language and Environment for Statistical Computing*. Vienna, Austria, 2013.
- [14] K. Radinsky and E. Horvitz. Mining the web to predict future events. In *WSDM*, 2013.
- [15] Telegraph.co.uk. [Twitter in numbers](#). Accessed: 2013-08-25.
- [16] TheGuardian.com. [John Kerry urges peace in Egypt amid anti-government protests](#). Accessed: 2013-08-25.
- [17] J. Ward. Hierarchical grouping to optimize an objective function. *J. Am. Stat. Assoc.*, 58(301):236–244, 1963.