

Actions Are Louder than Words in Social Media

Rostyslav Korolov^{*}, Justin Peabody[†], Allen Lavoie[†], Sanmay Das[†], Malik Magdon-Ismaïl[‡] and William Wallace^{*}

^{*}Department of Industrial and Systems Engineering, Rensselaer Polytechnic Institute
{korolr, wallaw}@rpi.edu

[†]Department of Computer Science & Engineering, Washington University in St. Louis
{peabody, allenlavoie, sanmay}@wustl.edu

[‡]Department of Computer Science, Rensselaer Polytechnic Institute
magdon@cs.rpi.edu

Abstract—We study the relationship between the level of chatter on a social medium (like Twitter) and the level of the observed actions related to the chatter. For example, in a disaster, how does relief-donation chatter on Twitter correlate with the dollar amount received? One hypothesis is that a fraction of those who act will also tweet about it, which implies linear scaling, $action \propto chatter$. On the other hand, if there is a contagion effect (those who tweet about donation incite others to donate) and these incited donors tend to be “quiet” and not broadcast their actions, then we expect superlinear scaling,

$$action \propto chatter^\gamma,$$

where $\gamma > 1$. We show, using a simple model, that the degree sequence of the social media “follower” network plays a key role in determining the scaling exponent γ . For random graphs and power-law graphs, the scaling exponent is at or near 2 (quadratic amplification). We empirically validate the model’s predictions using location-paired donation and social media data from U.S. states after Hurricane Sandy. Understanding the scaling behavior that relates social-media chatter to real physical actions is an important step for estimating the extent of a response and for determining social-media strategies to affect the response.

I. INTRODUCTION

The aftermath of a disaster typically sees a complex social and humanitarian response. A challenge that the agencies which direct the disaster response face is to predict the flow of donations from thousands of individuals and institutions in order to best plan the relief efforts [1]. The emergence of social media as a real-time flow of information about a large segment of the population could provide a solution. There is significant general interest in using social media to forecast significant societal events [2]. It could even be possible to affect the outcome using social influence and homophily to spread messages and ideas in social networks [3]. However, forecasting actions from social media presents unique challenges and necessitates the development and validation of new models.

We explicitly model the relationship between social media chatter and real-world actions. We develop a general theoretical model of social media with an (asymmetric) follower structure, like Twitter. We assume that the network consists of “loud” users, who broadcast their actions on social media, and “quiet” users, who are more passive in their use of social media (see for example [4]). We observe only the messages from loud users and our model relates the *observed* chatter to the expected actions under different models of influence and network structures. For example, if we were to observe a thousand tweets from New York related to donating in the

aftermath of Hurricane Sandy, and five hundred tweets from Pennsylvania, how would we expect actual donations from the two states to scale relative to each other?

We find that for several different models of both how influence propagates and how the graph is structured, the expected scaling between messages in social media and actions in the real world is superlinear, indicating that a simple proportional model would be flawed. In fact for several of these models, and under relatively standard assumptions, the prediction is that the scaling exponent should be close to quadratic. We test our model predictions using data gathered in the aftermath of Hurricane Sandy: a set of tweets relevant to the disaster; and a dataset of actual donation values categorized by state. We find that the square of the number of tweets from a given state related to donations in the aftermath of Hurricane Sandy is a superb predictor of actual donation amounts by state, achieving an R-squared of 0.9286 in an OLS regression. The quadratic variable is a significantly better predictor than linear or super-quadratic ones. We also show that the prediction is substantially better than can be achieved using the known demographic model of donation. This is convincing evidence that Twitter can be a far more accurate predictor of donations to be received than conventional techniques. While this paper focuses on monetary donations, it also demonstrates the general utility of this method, and indicates that Twitter would also be useful for forecasting in-kind donations, an issue particularly important for humanitarian logistics. We report theoretical and empirical evidence to support two hypotheses:

1. Observed actions are related to social media activity through a social network amplification. Our theoretical and empirical results suggest quadratic amplification.
2. Social media activity captures the social network influence on action in a way that demographic indicators cannot. Therefore, social media is invaluable for predicting actions.

A. Related Work

Our work is related to four major distinct streams of literature: (1) Work on analyzing the determinants of donation behavior; (2) Research on *crisis informatics*, which looks at the entire socio-technical ecosystem surrounding a crisis; (3) Work on the use of social media for forecasting; (4) Models of influence and cascade behavior in social networks.

There has been a great deal of research on who chooses to donate and why, both in a general context [5], [6] and, specifically, for disaster relief [7]. For example, Destro and

Holguin-Veras [8] study donations as a function of demographics in the aftermath of Hurricane Katrina. There is evidence that corporate donations are higher for firms located closer to a disaster site, possibly because the firms feel a greater sense of responsibility when the disaster is closer to their “home” [9]. One goal of understanding donation behavior is to better predict donation flows which helps relief agencies manage the recovery [1]. In contrast to general demographic indicators, social media offer a personalized and direct way of forecasting expected donations, and so a successful method based on social media would have great value in disaster response.

The new field of *crisis informatics* [10] studies the use of social media in crisis situations which includes work addressing information dissemination [11], [12], self-organization of volunteer responders during a crisis [13], and identifying the trustworthy messages on social media [14]. Our work explores a novel problem at the intersection of existing research on crisis informatics and machine learning: Can we use social chatter to forecast actions in the real world?

The idea of using information in social media to forecast important societal events has recently been explored in several domains, including Google’s “flu trends” project [15] and the US IARPA “Open Source Indicators” program (<http://www.iarpa.gov/index.php/research-programs/osi>) and has led to the development of systems like Embers [2]. Typical approaches in existing data-mining projects are model-free, whereas we take a more fundamental approach towards modeling the actual dynamics of interactions in the underlying social network. As far as we are aware, we are the first to address the problem of forecasting donation actions from social media messages.

Finally, our main contributions are deeply related to the literature on influence propagation and behavior in social networks. Traditionally, the literature on network cascades addresses how decisions made by nodes in a network affect the dynamics of the network itself [16], [17]. This has been studied primarily in the context of disease epidemics [18], product adoption [19], and social influence [3], [20]. There has also been research on optimal “seeding” of networks through a set of initial adopters [21], [22], [23]. Often the structure of the network (e.g. small world vs. preferential attachment vs. Erdős-Rényi) can make a significant difference in the outcomes [24]. There has been some interest in the effects of passivity on user influence [4], with many Twitter users having largely passive followings. Passivity can lead to a two-step flow of information [25], originating with a small number of users but often passing through intermediaries. Our work fits into the context of this literature, but is distinct in two significant ways. First we look at a different type of problem: how can we relate observed words to *completely unobserved* actions when only a fraction of users talk about their actions on social media? Second, at least initially we are interested in abstracting away from temporal aspects of influence. Instead, we have access to all broadcast messages over a short time-window as a whole, and want to predict the total number of donations based on these messages, rather than forecasting the flow of information through the network. Thus, a simple one-step model of broadcast influence propagation is sufficient for our purposes, although studying the temporal dynamics of influence in this context is a rich avenue for future work.

II. ACTION MODEL

Before we give the general action model, we give the intuition behind our main result:

The magnitude of the action A scales superlinearly in the magnitude of broadcasts B , with scaling exponent ≈ 2 .

Consider a setting with N users. A user is “loud” with probability ℓ , and “receptive” with probability r (a user can be both loud and receptive). A user is loud independently of whether she is receptive. A fraction b of loud users tweet about the action (which we observe), so the expected number of broadcasts about the action is $B = b\ell N$.

Each message from a loud user is received by all followers of that user (including the loud user himself). Suppose that each loud user (who tweets) has a following F proportional to N , so $F = \alpha N$, and that these follower-sets are disjoint. Then the total number of followers who receive a message is $b\ell N F = \alpha b\ell N^2$. A fraction r of these followers is receptive; assume that every receptive user that receives the broadcast will act. Then the number of actions is $A = r\alpha b\ell N^2$. Since $N^2 = B^2/b^2\ell^2$, we have that

$$A = \lambda B^2,$$

where $\lambda = r\alpha/b\ell$. It is helpful to isolate the key ingredients that lead to this result.

- i. Broadcasters have influence sets proportional to the number of broadcasters.
- ii. The influence sets are disjoint.
- iii. A fraction of those influenced will donate.

The result is robust to the specific details of the model, provided that these basic assumptions hold. The second assumption is interesting, and requires that the influence sets are not too large. When the influence sets are large, a receptive user may receive multiple broadcasts. If a receptive user acts with a probability proportional to the number of broadcast messages she gets then the result would still hold, but this may be slightly unrealistic due to diminishing returns with respect to the number of messages received. This would result in sub-quadratic scaling of D with B . We now turn to specifying our models and results more concretely.

A. Random Seed Model

Let the directed graph $G = (V, E)$ be the social media network with vertices representing actors and (directed) edges representing the follower relation. For an actor $v_i \in V$ let N_i be the neighborhood of actors that it is following; the (out-)degree of v_i is $\delta_i = |N_i|$. We allow self edges, which means that a node can “hear” itself. Let \mathcal{B} be a set of initial nodes selected independently and uniformly at random; the set \mathcal{B} contains the “loud” actors who tweet about the action. Let $B = |\mathcal{B}|$; $\mathbb{P}[v_i \in \mathcal{B}] = p$ and $\mathbb{E}[B] = Np$. The neighborhood of \mathcal{B} , denoted $\mathcal{N}(\mathcal{B})$, is the set of nodes which have links into the set \mathcal{B} . The set of nodes influenced by \mathcal{B} , denoted $\mathcal{I}(\mathcal{B})$ (when the context is clear just \mathcal{I}), is the set of nodes which are influenced by the initial broadcast messages started from \mathcal{B} . We consider three models of influence (or contagion):¹

¹The Contact and Excitation models are similar to the threshold and independent cascade models in influence propagation. The main difference is that in our process, the propagation stops after one step.

Contact:

Every node who hears a message to act is influenced; $\mathcal{I} = \mathcal{N}(\mathcal{B})$ (since we allow self-edges, a node can be in its own neighborhood and “hear” itself),

$$\mathbb{P}[v_i \in \mathcal{I}] = 1 - (1 - p)^{\delta_i}.$$

The formula follows because there are δ_i nodes in v_i 's neighborhood who can contact v_i ; each neighbor is independently put into \mathcal{B} with probability p ; and, v_i is influenced if *at least one* of its δ_i neighbors is in \mathcal{B} .

Excitation:

Every message a node hears has a probability α to excite the node into action. The node acts if any one of the messages excites it into action. Each potential influencer of v_i will excite v_i with probability αp . So,

$$\mathbb{P}[v_i \in \mathcal{I}] = 1 - (1 - \alpha p)^{\delta_i}.$$

This model is the contact model with p replaced by αp (with probability αp the neighbor is in the set of initial nodes, *and* sends a message to v_i that excites v_i into action. Though the process of the Excitation model is different from the Contact model, the mathematical form is similar with a simple change of the parameters. Thus, in the rest of the paper we only analyze the Contact model; it is also applicable to Excitation.

Proportionate:

A node is excited into action with a probability proportional to number of messages it receives. Let k_i be the number of nodes in v_i 's neighborhood that are broadcast seeds (which equals the number of messages v_i will get); k_i is a binomial random variable, $\mathbb{P}[k_i] = \binom{\delta}{k_i} p^{k_i} (1 - p)^{\delta - k_i}$, and $\mathbb{P}[v_i \in \mathcal{I} | k_i] = \beta k_i$. Therefore,

$$\begin{aligned} \mathbb{P}[v_i \in \mathcal{I}] &= \sum_{k_i=0}^{\delta_i} \mathbb{P}[v_i \in \mathcal{I} | k_i] \mathbb{P}[k_i] \\ &= \beta \sum_{k_i=0}^{\delta_i} k_i \binom{\delta_i}{k_i} p^{k_i} (1 - p)^{\delta_i - k_i} \\ &= \beta p \delta_i \end{aligned}$$

We assume that some fraction of the influence set acts, and so the number of actions $A \propto |\mathcal{I}|$. Thus, to understand how the number of actions observed depends on the number of messages B , one must compute the dependence of $|\mathcal{I}|$ on B .

One point worth noting here is that we are ignoring the *dynamics* of the message propagation since our goal is to predict aggregate behavior from aggregate observed chatter.

B. Analysis

We first analyze the Proportionate model. The expected magnitude of the action is $\mathbb{E}[A] = \sum_i \mathbb{P}[v_i \in \mathcal{I}]$ and so,

$$\mathbb{E}[A] = \beta p \sum_i \delta_i = \beta p |E|.$$

Since $\mathbb{E}[B] = p|V|$, the scaling behavior for A with respect to B is governed by how the density of the social network graph scales with $|V|$ (assuming p , the fraction of nodes that are broadcast seeds is a constant).

Erdős-Rényi Network Model. In the Erdős-Rényi random graph model for a social network, each directed edge exists with probability q (that could possibly depend on $|V|$).² Therefore, $|E| \propto |V|^{1+\gamma}$, where γ defines the sparsity of the graph, $0 \leq \gamma \leq 1$. The graph is dense if $\gamma = 1$ and sparse if $\gamma = 0$. Suppose that the initial seed probability $p = \rho|V|^{-\xi}$ where $0 < \xi \leq 1$ ($\xi = 1$ means a constant number of nodes is seeded; $\xi = 0$ means a constant fraction of nodes is seeded). We now get our first result.

Theorem 1. *For the Proportionate Action Model in an Erdős-Rényi social network with sparsity parameter γ and initial seed probability $p \propto |V|^{-\xi}$, the scaling law is $\mathbb{E}[A] \propto \mathbb{E}[B]^{1+\gamma/(1-\xi)}$.*

Proof: We saw above that $\mathbb{E}[A] \propto p|E| \propto |V|^{1-\xi+\gamma}$. Further, $\mathbb{E}[B] = p|V| \propto |V|^{1-\xi}$, so $|V| \propto \mathbb{E}[B]^{1/(1-\xi)}$. Hence $\mathbb{E}[A] \propto (\mathbb{E}[B]^{1/(1-\xi)})^{1-\xi+\gamma}$, which is the desired dependence. ■

For the special case of a constant edge probability, $\gamma = 1$ (dense graphs), and a constant fraction of nodes seeded, $\xi = 0$, we find that the scaling law is quadratic.

Power-Law Degree Distributions. A power-law degree distribution is specified by the number of nodes of degree 1 (k_1), and the decay rate $q > 2$. The power-law distribution is a very common assumption in social networks. It is known to be generated by a number of preferential attachment growth models such as the Price model [26] and the Barabasi-Albert model [27]. Typically $q \in (2, 3]$ in real networks, with the observed value of q often close to 2 in location-based social networks [28], [29]. The number of nodes with degree i is $k_i = k_1/i^q$. If the maximum degree is i_{max} , then

$$\begin{aligned} |V| &= \sum_{i=1}^{i_{max}} k_i = k_1 \sum_{i=1}^{i_{max}} \frac{1}{i^q} \approx k_1 \zeta(q); \\ |E| &= \sum_{i=1}^{i_{max}} i k_i = k_1 \sum_{i=1}^{i_{max}} \frac{1}{i^{q-1}} \approx k_1 \zeta(q-1), \\ &= |V| \frac{\zeta(q-1)}{\zeta(q)} \propto |V|. \end{aligned}$$

where the approximations above are in the asymptotic limit of very large social networks, when we can approximate the summations by taking the upper-limits to ∞ ($\zeta(\cdot)$ is the Riemann Zeta function). That is, in power-law graphs with constant power-law exponent $q > 2$ (independent of the number of nodes), the number of edges is asymptotically proportional to the number of nodes.

Theorem 2. *For the Proportionate Action Model in a power-law graph with parameters (k_1, q) and initial seed probability $p \propto |V|^{-\xi}$, the scaling law is $\mathbb{E}[A] \propto \mathbb{E}[B]^\gamma$, where*

$$\gamma = 1 + \frac{\log(\zeta(q-1)/\zeta(q))}{(1-\xi) \log k_1 \zeta(q)}.$$

Proof: We use $\mathbb{E}[A] \propto p|E| \propto |V|^{-\xi}|E|$, and $|V| \propto \mathbb{E}[B]^{1/(1-\xi)}$. We thus need to know how $|E|$ scales with $|V|$.

²For sparse graphs, $q = O(\frac{1}{|V|})$ and for dense graphs, $q = O(1)$.

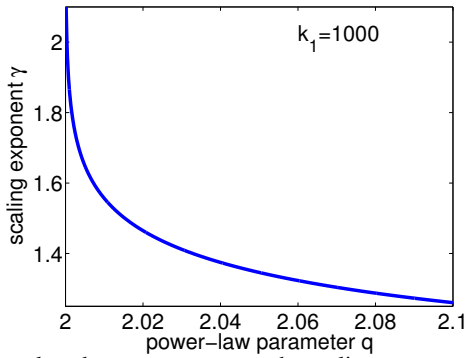


Fig. 1: Power-law decay parameter vs. the scaling exponent for $\xi = 0$.

Writing $|E| = |V|^\alpha$ we observe that

$$\alpha = \log |E| / \log |V| \approx \frac{\log k_1 + \log \zeta(q-1)}{\log k_1 + \log \zeta(q)},$$

and $\gamma = (\alpha - \xi) / (1 - \xi) = 1 + (\alpha - 1) / (1 - \xi)$. The result follows after substituting in α and some algebra. ■

As an example, for $k_1 = 1000$ (1,000 nodes of degree 1), we plot the scaling exponent γ versus the power-law decay parameter q for the case $\xi = 0$. We observe from Fig. 1 that when the power-law decay parameter is close to 2, the scaling exponent γ is in the vicinity of 2, that is quadratic scaling.

We now analyze the contact model. When the probability of a contact is small, then a node will typically have either zero or one contact with a broadcaster (more than one contact will be rare) and in this case the contact model and the Proportionate Action Model are similar. More generally, to get the scaling exponent, we need to analyze the size of the influenced set, $|Z| = |\mathcal{N}(B)|$. We have that

$$\mathbb{E}[|Z|] = \sum_i \mathbb{P}[v_i \in Z] = \sum_i (1 - (1-p)^{\delta_i})$$

We see that if p is small compared to $|V|$, which means a small fraction of the nodes are initially seeded (as is the case in practice), then we may approximate $1 - (1-p)^{\delta_i} \approx p\delta_i$, which reduces to the Proportionate model. Suppose that $p = \rho|V|^{-\xi}$ where $0 < \xi \leq 1$ ($\xi = 1$ means a constant number of nodes is seeded; $\xi = 0$ means a constant fraction of nodes is seeded). Then $\mathbb{E}[B] = \rho|V|^{1-\xi}$, and we can approximate

$$\mathbb{E}[A] \approx \rho|V|^{-\xi}|E|.$$

In this regime for p , the contact model is like the Proportionate model and asymptotically, we get the same scaling law exponents.

Theorem 3. *For the Contact model with seed probability $p \propto |V|^{-\xi}$ with $\xi > 0$, asymptotically in $|V|$, the scaling law is $\mathbb{E}[A] \propto \mathbb{E}[B]^\gamma$, where γ is the same scaling exponent obtained for the Proportionate model.*

So when the number of seeded nodes is $o(|V|)$, the Contact and Proportionate models are equivalent, independent of the actual social network model (Erdős-Rényi or power law).

We have shown that, for a variety of theoretical contagion models on different social network models, the scaling is superlinear, and the scaling exponent $\gamma \approx 2$. The scaling exponent is controlled almost entirely by the degree distribution of the social network. These results provide a theoretical basis




Fig. 2: Statewise data of \$-donated vs. tweet intensity (log-log)

for understanding the relationship between what is advertised on social media and the ultimate acts that ensue. For example, if we were estimating the extent of the Arab-spring demonstrations based on social media activity, our models would predict that if social media activity doubles, demonstration activity would (approximately) quadruple.

III. EXPERIMENTS

We validate our model on donation data from Hurricane Sandy, which hit the New York – New Jersey coast in October 2012. For this event, the social medium is Twitter and broadcast messages are donation-related. Actions are realized donations. We test our model using a cross-sectional analysis (stratified geographically) of the intensity of donation tweets on Twitter versus the intensity of donations received.

A. Data and Design

The dataset consists of 14,915,996 twitter messages from 10/25/2012 to 11/05/2012 containing at least one of the keywords or hashtags {'hurricane', 'Sandy', '#Hurricane', '#Sandy'}. These messages are extracted from the Twitter API which provides a random a sample of $\sim 1\%$ of all tweets

Extracting messages relevant to donation. We need to extract messages relating specifically to donation in the aftermath of Hurricane Sandy, disambiguating them from other messages also relevant to Sandy. We used the following procedure:

1. Create a training data set of 150 donation-relevant and 150 irrelevant messages (using a human to label the tweets).
2. Use a supervised algorithm with bag of words features to learn a classifier on the training data. We used an SVM implementation from the LIBSVM package [30].
3. Run the classifier on all the messages to identify the donation-relevant tweets.

Our procedure yielded 579,508 messages that the SVM classifier labeled as relevant to donations.

Geolocating Messages. To perform a cross-sectional analysis, we treat geographic regions, namely US states, as independent social networks. While there is certainly some cross-talk between states, we expect this to be a reasonable treatment. Twitter's public API only provides limited geographic information for messages, and this is only available if the user-settings permit. Therefore we need to *infer* message location based on properties of the message and the profile of author [31]. We used the software from [32] to infer locations. It was

possible to infer locations for 252,610 of the tweets selected as donation-relevant (roughly 50%) and 6,011,486 of all tweets in the dataset. Since the locations are not determined to the same level of granularity for each message, we only used messages for which the US County could be specified. Ultimately we have reduced the data set to 131,624 donation-relevant tweets for which we have county-level locations. It is important to note that none of this filtering should have disparate impact on donation-related messages, and in particular, that any non-linear relationship between donation-related tweets and donations should still hold, since non-locatable donation-related tweets should just be proportional to locatable ones.

Donation Data. To validate our model, we need actual donation data to compare with the intensity of tweet activity. We already have the tweets grouped by geolocation. The report on Hurricane Sandy response [33] gives us information on donations, grouped by US state. The report includes \$402,407,443 worth of donations from 624 corporate and institutional donors. Those include corporations (42% of value) and public and private charity institutions (the rest). 72% of corporate donations value is reported as coming from “giving programs”, which accumulate individual donations, so the majority of donation dollars represent accumulated individual donations, not large single corporate gifts. The report was published in October 2014 with data gathered prior to July 2014. We do not have dates for the donations themselves, and the report does not include all donations (e.g. in-kind gifts are omitted). Nevertheless, it is reasonable to assume that the donation data from this report is proportional to total amounts, at least for monetary donations.

Fitting the data. For states and territories $s = 1, \dots, 54$, we have donation-relevant tweet intensities b_1, \dots, b_{54} , where b_s is the number of tweets geolocated to within state s . We also have donation values a_1, \dots, a_{54} , where for state s , a_s is the donation value reported in [33]. We use ordinary least squares regression to study an explanatory linear model: $a_s = \alpha b_s^\gamma + \epsilon_s$, where ϵ_s is independent noise and γ is the scaling exponent.

B. Results

We test two hypotheses:

- H1** Donation amounts scale approximately quadratically in the number of donation-relevant Twitter messages.
- H2** Twitter activity is more informative about donation behavior than existing models that take into account the factors thought to most commonly affect donation behavior *a priori*, namely proximity, population, and income [8].

Hypothesis **H1** tests our theoretical model that relates *observed* social media behavior to *observed* actions. Hypothesis **H2** tests whether there is significant additional predictive power embedded in the social media activity as compared to traditional methods for predicting donation behavior.

H1 (The scaling law): To test the first hypothesis, we performed two tests. First, we found the best linear fit to a log-log plot of donation-relevant tweets and dollar value of donations at a state-by-state level (see Figure 2). The slope (which corresponds to the exponent) is 1.74, and the 99% confidence interval for the slope is [1.17, 2.31]. That is, with 99% confidence we can claim that the scaling exponent is




Fig. 3: R^2 (left) and LOO-Cross-Validation (right) for different scaling exponents γ : $a_s = \alpha \cdot b_s^\gamma + \epsilon_s$.

significantly larger than 1 (superlinear scaling) which is strong evidence in favor of a contagion effect in donation behavior. It is not simply the case that some fraction of donors tweet and the process stops (we would then see linear scaling); the tweets must incite more donations (contagion) *without associated further tweet activity* in order to get superlinear scaling. This indicates the presence of a substantial “quiet” population on twitter, who are receptive to influence and take action without tweeting. Further, the scaling is close to quadratic, as is predicted quantitatively by the theoretical model.

We also report R^2 goodness of fit and the leave-one-out cross validation error (LOO-CV) for different scaling exponents (Fig. 3). Clear peaks in predictive performance for $\gamma \approx 2$ corroborate our theoretical prediction.

Summary. Our results strongly support superlinearity in how donations scale with tweets - the presence of a social media amplification effect, as well as the specific quantitative predictions of our model which suggests the near-quadratic relation between words and actions observed.

H2 (Tweets vs. demographics predictors of donation): We compare the performance of a regression model that uses *only* the number of donation-relevant Twitter messages in a state against a regression that uses demographic factors known to be good predictors of donations: income, population and distance to disaster [8]. For the Twitter message regression, we use the scaling exponent $\gamma = 1.8$ suggested by the analysis of **H1**. The details of the regression are in the table below.

Coefficient	t -value	$R^2(R_{adj}^2)$	LOO-CV
0.845	26.3	93.01%(92.88%)	\$8.9M

We now compare with a regression on the important demographic factors identified in literature [8]: population, total income of residents, and distance from the disaster area for each state. The first two were taken from US Census Bureau data. The third (distance) is computed as the number of state borders that must be crossed to reach New York or New Jersey, emulating the proximity of a state to the disaster area. This model yields $R^2=55.30%$ ($R_{adj}^2=52.50%$) and leave-one-out L_2 cross-validation error of \$28.99M.

Summary. Our results indicate a surprising conclusion that emphasizes the importance of social media data. The regression which uses the *single* independent variable (donation tweet intensity) substantially outperforms the regression using three well known demographic independent variables. Square of the number of Twitter messages may be one of the best predictors of donation amounts presently available ($R^2 = 93%$ ($R_{adj}^2=92.88%$), whereas the benchmark regression only gives $R^2 = 53.3%$ ($R_{adj}^2=52.5%$)). Thus, Twitter messages are not merely a proxy for demographic information.

IV. DISCUSSION

We observe that social-media chatter is not just a proxy for population in forecasting real-world actions; significant value is added by using it to estimate the volume of action instead of the number of actors in the population. Even more, social media alone is significantly more powerful than several well known demographic factors that have been shown to predict donation. We gave a theoretical analysis for how social-chatter quantitatively relates to action via a superlinear scaling law (near-quadratic). This is due to the combination of contagion effects and the passivity or “quietness” of many social media users. Moreover, our model quantitatively predicts scaling exponents near 2. We validated our model on a particular event, donation behavior in the aftermath of Hurricane Sandy. The empirical data on tweets and dollar values of donations clearly supports a scaling exponent near 2.

Our model is a simple one-step model, and yet it gives results that are closely corroborated by the data. It would be interesting to extend our analysis to more complex multi-step diffusion models where an initial seed set broadcasts the message and a fraction of receptive nodes that receive the message and act might re-broadcast the message and so on. We also only considered a static snapshot of all chatter activity and all donations. It would be very interesting to understand if the dynamics of the chatter give a predictive edge in estimating the action. It would also be interesting to validate the model further on different types of contexts: evacuation; demonstration/congregation (strikes, flashmobs), etc.

ACKNOWLEDGMENT

The authors gratefully acknowledge support from NSF Grant 1124827. This material is also based partially upon work sponsored by the Army Research Lab under cooperative agreement number W911NF-09-2-0053 and by Department of Homeland Security through the Command, Control, and Interoperability Center for Advanced Data Analysis Center of Excellence. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory, or the U.S. Government.

REFERENCES

- [1] J. Holguín-Veras, M. Jaller, L. N. Van Wassenhove, N. Pérez, and T. Wachtendorf, “On the unique features of post-disaster humanitarian logistics,” *J. Operations Management*, vol. 30, no. 7, pp. 494–506, 2012.
- [2] A. Doyle, G. Katz, K. Summers, C. Ackermann, I. Zavorin, Z. Lim, S. Muthiah, P. Butler, N. Self, L. Zhao, C.-T. Lu, R. P. Khandpur, Y. Fayed, and N. Ramakrishnan, “Forecasting significant societal events using the Embers streaming predictive analytics system,” *Big Data*, vol. 2, no. 4, pp. 185–195, 2014.
- [3] D. Cosley, D. P. Huttenlocher, J. M. Kleinberg, X. Lan, and S. Suri, “Sequential influence models in social networks.” *AAAI Int. Conf. on Weblogs and Social Media*, vol. 10, pp. 26–33, 2010.
- [4] D. M. Romero, W. Galuba, S. Asur, and B. A. Huberman, “Influence and passivity in social media,” in *Proc. ECML/PKDD*, 2011, pp. 18–33.
- [5] Y.-K. Lee and C.-T. Chang, “Who gives what to charity? Characteristics affecting donation behavior,” *Social Behavior and Personality*, vol. 35, no. 9, pp. 1173–1180, 2007.
- [6] R. Bekkers and P. Wiepking, “A literature review of empirical studies of philanthropy: Eight mechanisms that drive charitable giving,” *Nonprofit and Voluntary Sector Quarterly*, 2010.
- [7] C. Torrey, M. Burke, M. Lee, A. Dey, S. Fussell, and S. Kiesler, “Connected giving: Ordinary people coordinating disaster relief on the Internet,” in *Hawaii International Conference on System Sciences*, 2007.
- [8] L. Destro and J. Holguín-Veras, “Material convergence and its determinants,” *Transportation Rsrch. Rec.*, vol. 2234, no. 1, pp. 14–21, 2011.
- [9] A. Muller and G. Whiteman, “Exploring the geography of corporate philanthropic disaster response: A study of Fortune Global 500 firms,” *J. Business Ethics*, vol. 84, no. 4, pp. 589–603, 2009.
- [10] L. Palen and S. B. Liu, “Citizen communications in crisis: Anticipating a future of ICT-supported public participation,” in *Proc. SIGCHI*. ACM, 2007, pp. 727–736.
- [11] K. Starbird and L. Palen, “(How) will the revolution be retweeted? Information diffusion and the 2011 Egyptian uprising,” in *Proc. ACM CSCW*. ACM, 2012, pp. 7–16.
- [12] Y. Tyshchuk, H. Li, H. Ji, and W. A. Wallace, “Evolution of communities on twitter and the role of their leaders during emergencies,” in *Proc. ASONAM*, pp. 727–733.
- [13] K. Starbird and L. Palen, “Voluntweeters: Self-organizing by digital volunteers in times of crisis,” in *Proc. SIGCHI*, 2011, pp. 1071–1080.
- [14] C. Hagar, “Crisis informatics: Perspectives of trust-is social media a mixed blessing?” *SLIS Student Research J.*, vol. 2, no. 2, p. 2, 2013.
- [15] S. Cook, C. Conrad, A. L. Fowlkes, and M. H. Mohebbi, “Assessing Google flu trends performance in the United States during the 2009 Influenza virus A (H1N1) pandemic,” *PLoS one*, vol. 6, no. 8, 2011.
- [16] W. B. Arthur, “Competing technologies, increasing returns, and lock-in by historical events,” *The Economic Journal*, pp. 116–131, 1989.
- [17] S. Bikhchandani, D. Hirshleifer, and I. Welch, “Learning from the behavior of others: Conformity, fads, and informational cascades,” *The Journal of Economic Perspectives*, pp. 151–170, 1998.
- [18] S. Eubank, H. Guclu, V. A. Kumar, M. V. Marathe, A. Srinivasan, Z. Toroczkai, and N. Wang, “Modelling disease outbreaks in realistic urban social networks,” *Nature*, vol. 429, no. 6988, pp. 180–184, 2004.
- [19] J. Leskovec, A. Singh, and J. Kleinberg, “Patterns of influence in a recommendation network,” in *Advances in KDD*, 2006, pp. 380–389.
- [20] S. Adali, F. Sisenda, and M. Magdon-Ismael, “Actions speak as loud as words: Predicting relationships from social behavior data,” in *Proc. WWW*, 2012, pp. 689–698.
- [21] P. Domingos and M. Richardson, “Mining the network value of customers,” in *Proc. KDD*, 2001, pp. 57–66.
- [22] D. Kempe, J. Kleinberg, and É. Tardos, “Maximizing the spread of influence through a social network,” in *Proc. KDD*, 2003, pp. 137–146.
- [23] E. Anshelevich, A. Hate, and M. Magdon-Ismael, “Seeding influential nodes in non-submodular models of information diffusion,” in *Proc. AAMAS*, 2013, pp. 1305–1306.
- [24] D. J. Watts, “A simple model of global cascades on random networks,” *PNAS*, vol. 99, no. 9, pp. 5766–5771, 2002.
- [25] S. Wu, J. M. Hofman, W. A. Mason, and D. J. Watts, “Who says what to whom on Twitter,” in *Proc. WWW*. ACM, 2011, pp. 705–714.
- [26] D. d. S. Price, “A general theory of bibliometric and other cumulative advantage processes,” *JASIST*, vol. 27, no. 5, pp. 292–306, 1976.
- [27] A.-L. Barabási and R. Albert, “Emergence of scaling in random networks,” *Science*, vol. 286, no. 5439, pp. 509–512, 1999.
- [28] N. Li and G. Chen, “Analysis of a location-based social network,” in *IEEE Intl. Conf. on Comp. Sci. and Engg.*, vol. 4, 2009, pp. 263–270.
- [29] S. Scellato and C. Mascolo, “Measuring user activity on an online location-based social network,” in *INFOCOM Workshops*. IEEE, 2011, pp. 918–923.
- [30] C.-C. Chang and C.-J. Lin, “LIBSVM: A library for Support Vector Machines,” *ACM TIST*, vol. 2, no. 3, pp. 27.1–27.27, 2011.
- [31] S. Kumar, F. Morstatter, and H. Liu, *Twitter Data Analytics*. New York, NY, USA: Springer, 2014.
- [32] M. Dredze, M. Paul, S. Bergsma, and H. Tran, “Carmen: A Twitter geolocation system with applications to public health,” in *HIAI Workshop (at AAAI’13)*, 2013, pp. 20–24.
- [33] The Foundation Center, “Philanthropy & Hurricane Sandy. A Report on the Foundation and Corporate Response.” 2014. [Online]. Available: http://foundationcenter.org/gainknowledge/research/pdf/sandy_philanthropy_2014.pdf