Who Will Retweet This? Detecting Strangers from Twitter to Retweet Information

KYUMIN LEE, Utah State University
JALAL MAHMUD, IBM Research – Almaden
JILIN CHEN, Google
MICHELLE ZHOU, IBM Research – Almaden
JEFFREY NICHOLS, Google

31

There has been much effort on studying how social media sites, such as Twitter, help propagate information in different situations, including spreading alerts and SOS messages in an emergency. However, existing work has not addressed how to actively identify and engage the right strangers at the right time on social media to help effectively propagate intended information within a desired time frame. To address this problem, we have developed three models: (1) a feature-based model that leverages people's exhibited social behavior, including the content of their tweets and social interactions, to characterize their willingness and readiness to propagate information on Twitter via the act of retweeting; (2) a wait-time model based on a user's previous retweeting wait times to predict his or her next retweeting time when asked; and (3) a subset selection model that automatically selects a subset of people from a set of available people using probabilities predicted by the feature-based model and maximizes retweeting rate. Based on these three models, we build a recommender system that predicts the likelihood of a stranger to retweet information when asked, within a specific time window, and recommends the top-N qualified strangers to engage with. Our experiments, including live studies in the real world, demonstrate the effectiveness of our work.

Categories and Subject Descriptors: H.5.2 [**Information Interfaces and Presentation**]: User Interfaces

General Terms: Design, Algorithms, Experimentation

Additional Key Words and Phrases: Twitter, retweet, social media, willingness, personality

ACM Reference Format:
Kyumin Lee, Jalal Mahmud, Jilin Chen, Michelle Zhou, and Jeffrey Nichols, 2015. Who will retweet this? detecting strangers from twitter to retweet information. ACM Trans. Intell. Syst. Technol. 6, 3, Article 31 (April 2015), 25 pages.
DOI: http://dx.doi.org/10.1145/2700466

1. INTRODUCTION

With the widespread use of social media sites like Twitter and Facebook and the ever-growing number of users, there has been much effort on understanding and modeling information propagation on social media [Agarwal et al. 2008; Bakshy et al. 2011; Cha et al. 2010; Goyal et al. 2010; Huang et al. 2012; Romero et al. 2011; Singer 2012; Ver Steeg and Galstyan 2012; Weng et al. 2010].

An early version of this article appeared in the 2014 ACM Proceedings of the IUI conference [Lee et al. 2014].
Authors' addresses: K. Lee, Department of Computer Science, 4205 Old Main Hill Logan, UT 84322-4205; email: kyumin.lee@usu.edu; J. Mahmud and M. Zhou, 650 Harry Road, San Jose, CA 95120; emails: {jumahmud, mzhou}@us.ibm.com; J. Chen and J. Nichols, 1600 Amphitheatre Pkwy, Mountain View, CA 94043; emails: jilinc@acm.org, jeff@jeffreynichols.com. The work was done while J. Chen and J. Nichols were at IBM.

DOI: http://dx.doi.org/10.1145/2700466

Most of the work assumes that information is propagated by a small number of influential volunteers who possess certain qualities, such as having a large number of followers, that make them extremely effective in propagating information [Starbird and Palen 2010]. For example, these users can help spread emergency alerts, such as fire hazards or SOS messages like requesting blood donations, to reach more people faster.

However, prior research efforts ignore several critical factors in influencer-driven information propagation. First, influential users may be unwilling to help propagate the intended information for various reasons. For example, they may not know the truthfulness of a piece of information, and thus are unwilling to risk their reputation to spread the information. Second, an influential user may be unavailable to help propagate information when needed. For example, influential users may not be online to help propagate SOS messages when a disaster strikes.

Since everyone is potentially an influencer on social media and is capable of spreading information [Bakshy et al. 2011], our work aims to identify and engage the right people at the right time on social media to help propagate information when needed. We refer to these people as *information propagators*. Since not everyone on social media is willing or ready to help propagate information, our goal is to model the characteristics of information propagators based on their social media behavior. We can then use the established model to predict the likelihood of a person on social media being an information propagator. As the first step, we focus on modeling *domain-independent* traits of information propagators, specifically their *willingness* and *readiness* to spread information.

In many situations, including emergency or disastrous situations, information propagation must be done within a certain time frame to optimize its effect. To satisfy such a time constraint, we thus also develop a wait-time model based on a user's previous retweeting wait times to predict the user's next retweeting time when asked.

For the sake of concreteness, in this article, we focus on Twitter users, although our core technology can be easily applied to other social media platforms. On Twitter, the most common method for propagating information is retweeting,[1] which is to repost others' tweets in your own content stream. Our work is thus reduced to the problem of finding strangers on Twitter who will retweet a message when asked.

To model one's willingness and readiness to retweet information, we first identify a rich set of features to characterize the candidate, including derived personality traits, social network information, social media activity, and previous retweeting behavior. Unlike existing work, which often uses only social network properties, our feature set includes *personality traits* that may influence one's retweeting behavior. For example, when asked by a stranger in an emergency, a person with a high level of altruism may be more responsive and willing to retweet. Similarly, a more active user who frequently posts status updates or reposts others' tweets may be more likely to retweet when asked. Our features capture a variety of characteristics that are likely to influence one's retweeting behavior.

To predict one's likelihood to retweet when asked, we train statistical models to infer the weights of each feature, which are then used to predict one's likelihood to retweet. Based on the prediction models, we also build a real-time recommender system that can rank and recommend the top-N candidates (*retweeters*) to engage with on Twitter.

To demonstrate the effectiveness of our work, we have conducted extensive experiments, including live studies in the real world. Compared to two baselines, our approach

[1]We use the terms "repost," "retweet," and "propagate" interchangeably.

significantly improves the *retweeting rate*:[2] the ratio between the number of people who retweeted and the number of people asked. To the best of our knowledge, our work is the first to address how to *actively* identify and engage strangers on Twitter to help retweet information. As a result, our work offers four unique contributions:

- A feature-based model including one's personality traits for predicting the likelihood of a stranger on Twitter to retweet a particular message when asked.
- A wait-time model based on a person's previous retweeting wait times to estimate his or her next retweeting wait time when asked.
- A subset selection model that *automatically* selects a subset of people from a set of available people using probabilities predicted by the feature-based model and maximizes retweeting rate.
- A retweeter recommender system that uses the three models mentioned previously to effectively select the right set of strangers on Twitter to engage with in real time.

2. RELATED WORK

Our work is most closely related to the recent efforts on actively engaging strangers on social media for accomplishing certain tasks [Mahmud et al. 2013; Nichols and Kang 2012]. However, ours is the first on modeling and engaging strangers on social media to aid information propagation within a given time window.

Our work is also related to the effort on characterizing retweeters and their retweeting behavior [Macskassy and Michelson 2011]. However, the existing work does not include personality features as our model does. More importantly, unlike the existing model focusing on *voluntary* retweeting behavior, ours examines a person's retweeting behavior at the request of a stranger.

Some researchers studied recommending personalized tweets. For example, Chen et al. [2012] make use of different information to help recommendation, including the user's own tweet history, retweet history, and social relations between users. Their method of tweet recommendation makes use of collaborative ranking to capture personal interests. Feng and Wang [2013] developed a predictive model to rank the tweets according to their probability of being retweeted. Such ranked list of tweets is recommended to the user for retweeting. In contrast to the work that recommends tweets from users' followees on Twitter, we focus on recommending retweeters. Some of our features such as personality, past retweeting rate, and readiness-based features are not used in these papers. Furthermore, our work also addresses different objectives such as maximizing the retweeting rate or information reach, which were not the focus of the previous works.

There are many efforts on modeling influential behavior in social media. Such work finds influential users by their social network properties [Bakshy et al. 2011; Cha et al. 2010; Goyal et al. 2010; Huang et al. 2012; Singer 2012; Weng et al. 2010], content of posts [Agarwal et al. 2008], information forwarding/propagating activity [Romero et al. 2011], and information flow [Ver Steeg and Galstyan 2012]. In comparison, our work focuses on individuals' characteristics that influence their willingness and readiness to retweet at a stranger's request. Some of these characteristics, such as personality and readiness to retweet, have not been studied before.

As our goal is to support effective information diffusion, our work is related to efforts in this space. Bakshy et al. [2012] examine the role of the social network and the effects of tie strength in information diffusion. Hodas and Lerman [2014] show that the position of exposing messages on the user interface strongly affects social contagion.

[2]We use the terms "information propagation rate," "information repost rate," and "retweeting rate" interchangeably.

Fig. 1. An example Twitter account created for "public safety" data collection.

Chaoji et al. [2012] show how to maximize content propagation in one's own social network. In contrast, our approach aims at selecting a right set of *strangers* on social media to help spread information. Budak et al. [2011] have studied a different type of information diffusion, which spreads messages to counter malicious influences and hence minimize the influence of such campaigns. They proposed to identify a subset of individuals to start a counter campaign based on a set of viral diffusion features, including user virality and susceptibility and item virality [Hoang and Lim 2012]. These features are complementary to the features that we use, such as personality, messaging activity, and past retweeting activity. Moreover, there is little work on automatically identifying and engaging the *right* strangers at the *right* time on social media to aid information propagation as ours does.

3. CREATING GROUND-TRUTH DATASETS

Since there is no publicly available ground-truth data with which we can train and build our predictive models, we collected two real-world datasets. We created a total of 17 Twitter accounts, and our system automatically sent retweeting requests to 3,761 strangers on Twitter. Our first dataset examines *location-based targeting*, where people who live in a particular location were asked to retweet information relevant to that location. The second examines *topic-based targeting*, where people interested in a certain topic were asked to retweet information relevant to that topic.

We hypothesize that information relevance influences a person's retweeting behavior, especially at the request of a stranger. For example, people might be more likely to retweet news about public safety in an area where they live or work rather than for other locations. Similarly, a person might be more willing to retweet information on a topic in which he or she is interested.

Our dataset for location-based targeting (named "public safety") and the dataset for topic-based targeting (named "bird flu") are intended to examine how different types of information (location vs. topic) may impact retweeting behavior.

Public Safety Data Collection: For location-based targeting, we chose the San Francisco Bay Area as the location and sent tweets about local public safety news to people whom we identified as living or staying in that area. First, we created nine accounts on Twitter. All accounts had the same profile name, "Public Safety News," and the same description (Figure 1).

Note that we created multiple accounts to send a few messages per hour from each account in order to create a reasonable pretense of human behavior. Furthermore, previous studies have shown that if we were not careful, target strangers would silently flag an account as spam to cause the suspension of the account by Twitter [Mahmud et al. 2013; Nichols and Kang 2012]. Creating multiple accounts helped us avoid this possibility, and thus increased the number of users that we could reasonably contact per hour (each user received only one message).

Creating multiple accounts for research purposes is a commonly used methodology [Lee et al. 2010; Lee et al. 2011]. To make these accounts appear to be genuine, all accounts followed four to 10 users and had 19 followers. We also created the following and follower accounts, and some were also followed by the original accounts. We posted 11 public safety messages using each of the nine accounts before we contacted anyone on Twitter. We identified 34,920 Bay Area Twitter users using the Twitter Streaming API[3] with a geolocation filter corresponding to the Bay Area in June 2012. This stream retrieved only tweets that were marked as being sent within a bounding box equivalent to the Bay Area determined by using the Google Geocoding API.[4] We filtered out non-English tweets in this stream and created a list of unique users whose tweets were in the stream.

Among all the identified Twitter users, we randomly selected 1,902 people. From our public safety accounts, our system automatically sent messages to those people using the Twitter API and ensured that each person received only one message to avoid overburdening the person. Here is an example message sent:

> *@ SFtargetuser "A man was killed and three others were wounded in a shooting... http://bit.ly/KOl2sC" Plz RT this safety news"*

Each message contained the target person's screen name, the title of a news article obtained from a local news media site, a link to the article, and a phrase asking the person to retweet the message. The original link URL was shortened with the bit.ly URL shortening service to allow us to track user clicks on the link. Per our request, 52 of the 1,902 (2.8%) people retweeted our message, which reached a total of 18,670 followers of theirs.

Bird Flu Data Collection: For topic-based targeting, we chose people who tweeted about "bird flu," a topic commonly being discussed at the time of our study. First, we created eight accounts on Twitter (Figure 2). All accounts followed two to five users and had 19 followers. The following and followers accounts were created using the same method as in the public safety scenario. We then collected 13,110 people's profiles using the Twitter Search API and the queries "bird flu," "H5N1," and "avian influenza" in June 2012. We excluded non-English tweets and randomly selected 1,859 users. A message was then automatically sent to each selected person. Here is an example message sent:

> *@birdflutargetuser Plz RT bird flu news "Bird Flu viruses could evolve in nature http://bit.ly/MQBASY"*

As in the public safety study, the news articles were obtained from the news media sites. Out of the 1,859 users, 155 (8.4%) retweeted our messages, which reached their 184,325 followers.

For both datasets, through the Twitter API we collected publicly available information of each person whom we asked to retweet. This included their profile, people they

---

[3]http://dev.twitter.com/pages/streaming_api.

[4]https://developers.google.com/maps/documentation/geocoding/.

Fig. 2. An example Twitter account created for "bird flu" data collection.

followed, their followers, up to 200 of their most recently posted messages, and whether they retweeted our message (the ground truth).

4. FEATURE EXTRACTION

To model a person's likelihood to retweet, we have identified six categories of features, as described next.

4.1. Profile Features

Profile features are extracted from a user's Twitter profile and consist of *longevity (age) of an account, length of screen name, whether the user profile has a description, length of the description, and whether the user profile has a URL*. Our hypothesis behind the use of these features is that a user with a richer profile or a longer account history may be more knowledgeable in using advanced social media features, such as retweeting. Hence, when asked, they are more likely to retweet than those who have just opened an account recently or have little information in their profile.

4.2. Social Network Features

We use the following features to characterize a user's social network: *number of users following (friends)*, *number of followers*, and the *ratio of number of friends to number of followers*. These features indicate the "socialness" of a person. Intuitively, the more social a person is (e.g., a good number of followers), the more likely the person may be willing to retweet. These features may also signal potential motivations for retweeting (e.g., an act of friendship and to gain followers) [Boyd et al. 2010]. However, a person (e.g., a celebrity) with an extraordinary number of followers may be unwilling to retweet per a stranger's request.

4.3. Personality Features

Researchers have found that word usage in one's writings, such as blogs and essays, are related to one's personality [Fast and Funder 2008; Gill et al. 2009; Pennebaker et al. 2001]. In particular, Linguistic Inquiry and Word Count (LIWC) is used to analyze text statistically and find psychologically meaningful categories [Pennebaker et al. 2001]. Inspired by the existing work, we used the LIWC-2001 dictionary to compute

Table I. Correlations of a Big5 Facet with LIWC Categories

Top 20 Correlations of Friendliness with LIWC categories
Friends (0.23), Leisure (0.22), 1st Person Pl. (0.22), Family (0.2), Other Refs. (0.18), Up (0.18), Social Processes (0.17), Positive Emotions (0.17), Sexual (0.16), Space (0.16), Physical States (0.15), Home (0.15), Sports (0.15), Motion (0.14), Music (0.14), Inclusive (0.14), Eating (0.14), Time (0.13), Optimism (0.13), Causation (−0.13)

one's personality features. LIWC-2001 defines 68 different categories, each of which contains several dozen to hundreds of words. For each person, we computed his or her LIWC-based personality feature in each category as follows:

> Let g be an LIWC category, N_g denotes the number of occurrences of words in that category in one's tweets, and N denotes the total number of words in his or her tweets. A score for category g is then N_g/N.

Psychologists have developed several models of human personality. One of the more accepted models is the Big5 framework of personality traits [Costa and McCrae 1992], which proposes five key traits: *neuroticism*, *extraversion*, *openness*, *agreeableness*, and *conscientiousness*. Previous works, such as Fast and Funder [2008] and Gill et al. [2009], reveal correlations between the Big5 personality traits and the LIWC-category-based features extracted from text, such as blogs and essays. More recently, Yarkoni [2010] showed that correlations exist among LIWC features and lower-level facets of Big5. Motivated by the previous findings, we used the Big5 lower-level facets as well as the Big5 traits themselves as additional personality features in our model.

All previous works use the results of personality tests taken by their participants to determine the values of the Big5 features. However, their approach requires that users take a personality test, which is not practical in our situation. To derive personality scores for each of the Big5 dimensions and their lower-level facets, we use an alternative approach. We use the coefficients of correlation between Big5 lower-level facets and LIWC categories found by Yarkoni to compute those facet-level feature values. For example, Table I shows example correlations for a Big5 facet feature. To derive a feature value for a lower-level facet, we use a linear combination of LIWC categories (for which correlation was found statistically significant by Yarkoni), where correlation coefficients are used as weights. Yarkoni also reports correlation values between LIWC category features and Big5 traits. We use such correlations as weights for deriving Big5 feature values from LIWC-category-level features.

Overall, we computed 103 personality features from one's tweets: 68 LIWC features (e.g., word categories such as "sadness") and five Big5 dimensions (e.g., *agreeableness* and *conscientiousness*) with their 30 subdimensions. These features may signal potential motivations for retweeting (e.g., an act of altruism and to gain followers) [Boyd et al. 2010].

4.4. Activity Features

This feature category captures people's social activities. Similar to the reasons stated earlier, our hypothesis is that the more active people are, the more likely they would retweet when asked by a stranger. Moreover, new Twitter users or those who rarely tweet may not be familiar with the retweeting feature and be less likely to reweet. To evaluate this hypothesis, we use the following features:

- *Number of status messages*
- *Number of direct mentions (e.g., @johny) per status message*
- *Number of URLs per status message*
- *Number of hashtags per status message*

Table II. Readiness Features and Their Computations

Readiness Features	Computation
Tweeting likelihood of the day	T_D/N, where T_D is the number of tweets sent by the user on day D and N is the total number of tweets
Tweeting likelihood of the hour	T_H/N, where T_H is the number of tweets sent by the user on hour H and N is the total number of tweets
Tweeting likelihood of the day (entropy)	Entropy of tweeting likelihood of the day (T_D/N)
Tweeting likelihood of the hour (entropy)	Entropy of tweeting likelihood of the hour (T_H/N)
Tweeting steadiness	$1/\sigma$, where σ is the standard deviation of the elapsed time between consecutive tweets of users, computed from users' most recent K tweets (where K is set, e.g., to 20)
Tweeting inactivity	$T_R - T_L$, where T_R is the time the request was sent and T_L is the time the user last tweeted

- *Number of status messages per day during a person's entire account life (= total number of posted status messages / longevity)*
- *Number of status messages per day during last 1 month*
- *Number of direct mentions per day during last 1 month*
- *Number of URLs per day during last 1 month*
- *Number of hashtags per day during last 1 month*

These features also help us distinguish "sporadic" versus "steady" activeness. We hypothesize that "steady" users are more dependable and are more likely to retweet when asked. For each person, we computed these features based on their 200 most recent tweets, as our experiments have shown that 200 tweets are a good representative sample for deriving one's features.

4.5. Past Retweeting Features

We capture retweeting behavior with these features:

- *Number of retweets per status message: R/N*
- *Average number of retweets per day*
- *Fraction of retweets for which original messages are posted by strangers who are not in a person's social network*

Here R is the total number of retweets and N is the total number of status messages. We hypothesize that frequent retweeters are more likely to retweet in the future. The third feature measures how often a person retweets a message originated outside of the person's social network. We hypothesize that people who have done so are more likely to retweet per a stranger's request to do so.

4.6. Readiness Features

Even if a person is willing to retweet per a request, he or she may not be ready to do so at the time of the request for various reasons, such as being busy or not being connected to the Internet. Since such a context could be quite diverse, it is difficult to model one's readiness precisely. We thus use the following features (listed in Table II) to approximate readiness based on one's previous activity:

- *Tweeting likelihood of the day*
- *Tweeting likelihood of the hour*
- *Tweeting likelihood of the day (entropy)*
- *Tweeting likelihood of the hour (entropy)*
- *Tweeting steadiness*
- *Tweeting inactivity*

The first two features are computed as the ratio of the number of tweets sent by the person on a given day/hour and the total number of tweets. The third and fourth features measure entropy of tweeting likelihood of the day and the hour, respectively [Shannon 1948]. To follow is a person's (u) entropy of tweeting likelihood of the hour $P(x_1), P(x_2), P(x_3) \ldots P(x_n)$:

$$Entropy\,(u) = -\sum_{i=1}^{n} P\,(x_i) \log P(x_i).$$

In the previous equation, n is 24 to estimate the daily likelihood to tweet. The tweeting steadiness feature is computed as $1/\sigma$, where σ is the standard deviation of the elapsed time between consecutive tweets, computed from the most recent K tweets (where K is set to 20). The tweeting inactivity feature is the difference between the time when a retweeting request is sent and the time when the user last tweeted.

Our rationale for choosing this set of features is twofold. First, these features are good indicators of one's readiness from a particular aspect. For example, the value of *tweeting inactivity* may hint at one's availability, as a larger value may indicate either that the person is busy and hence uninterruptible or that he or she is out of reach. Second, these features are easy and fast to compute based on one's past tweeting activity instead of the tweet content.

5. PREDICTING RETWEETERS

Based on the features described earlier, we train a model to predict a user's likelihood to be a retweeter.

Training and Testing Sets. First, we randomly split each dataset (public safety and bird flu) into training (containing two-thirds of the data) and testing sets (containing one-third of the data). The two sets were stratified and contained the same ratio of retweeters and nonretweeters. Finally, for public safety, the training set had 35 retweeters and 1,233 nonretweeters, and the testing set had 17 retweeters and 617 nonretweeters. For bird flu, the training set had 103 retweeters and 1,136 nonretweeters, and the testing set had 52 retweeters and 568 nonretweeters. For each person in the sets, we computed all the features described previously.

Predictive Models. We compared the performance of five popular models: Random Forest, Naïve Bayes, Logistic Regression, SMO (SVM), and AdaboostM1 (with Random Forest as the base learner). We used WEKA [Hall et al. 2009] implementation of these algorithms and trained these models to predict the probability of a person to retweet and classify a person as a retweeter or nonretweeter. If a person's retweeting probability is greater than 0.5, the person will be classified as a retweeter.

Handling Class Imbalance. Both our datasets have an imbalanced class distribution: only 52 out of 1,902 users (2.8%) in the public safety dataset and 155 out of 1,859 users (8.4%) in the bird flu dataset were retweeters. Imbalanced class distribution in a training set hinders the learning of representative sample instances, especially the minority class instances, and prevents a model from correctly predicting an instance label in a testing set. The class imbalance problem has appeared in a large number of domains, such as medical diagnosis and fraud detection. There are several approaches to the problem, including oversampling minority class instances, undersampling majority class instances, and adjusting the weights of instances. Currently, we used both oversampling and weighting approaches to our class imbalance problem. For oversampling, we used the SMOTE [Chawla et al. 2002] algorithm. For weighting, we used a

Table III. 21 Features Selected by χ^2 in Public Safety Dataset

Feature Group	Public Safety Features
Profile	longevity of the account
Social network	|following| ratio of number of friends to number of followers
Activity	**|URLs| per day** **|direct mentions| per day** **|hashtags| per day** |status messages| |status messages| per day during entire account life |status messages| per day during last one month
Past retweeting	**|retweets| per status message** **|retweets| per day**
Readiness	Tweeting likelihood of the day Tweeting likelihood of the day (entropy)
Personality	7 LIWC features: **Inclusive**, Achievement, Humans, Time, Sadness, Articles, Nonfluencies 1 Facet feature: Modesty

Table IV. 46 Features Selected by χ^2 in Bird Flu Dataset

Feature Group	Bird Flu Features
Profile	length of description has description in profile
Activity	**|URLs| per day** **|direct mentions| per day** **|hashtags| per day** |URLs| per status message |direct mentions| per status message |hashtags| per status message
Past Retweeting	**|retweets| per status message** **|retweets| per day** |URLs| per retweet message
Readiness	Tweeting Likelihood of the Hour (Entropy)
Personality	34 LIWC features: **Inclusive**, Total Pronouns, 1st Person Plural, 2nd Person, 3rd Person, Social Processes, Positive Emotions, Numbers, Other References, Occupation, Affect, School, Anxiety, Hearing, Certainty, Sensory Processes, Death, Body States, Positive Feelings, Leisure, Optimism, Negation, Physical States, Communication 8 Facet features: Liberalism, Assertiveness, Achievement Striving, Self-Discipline, Gregariousness, Cheerfulness, Activity Level, Intellect 2 Big5 features: Conscientiousness, Openness

cost-sensitive approach of adding more weight to the minority class instances [Liu and Zhou 2006].

Feature Analysis. To improve the performance of our models, we analyzed the significance of our features using the training set. We computed the χ^2 value for each feature to determine its discriminative power [Yang and Pedersen 1997] and eliminated the features that do not contribute significantly to the result. Our analyses found 21 and 46 significant features for the two datasets, respectively (Tables III and IV). Note that these features consistently had positive discriminative power under a threefold cross-validation setting.

Several feature groups have more significant power distinguishing between retweeters and nonretweeters: *activity*, *personality*, *readiness*, and *past retweeting*. Although our two datasets are quite different, we found six significant features common to both sets (bolded in Tables III and IV). This suggests that it is possible to build *domain-independent* models to predict retweeters.

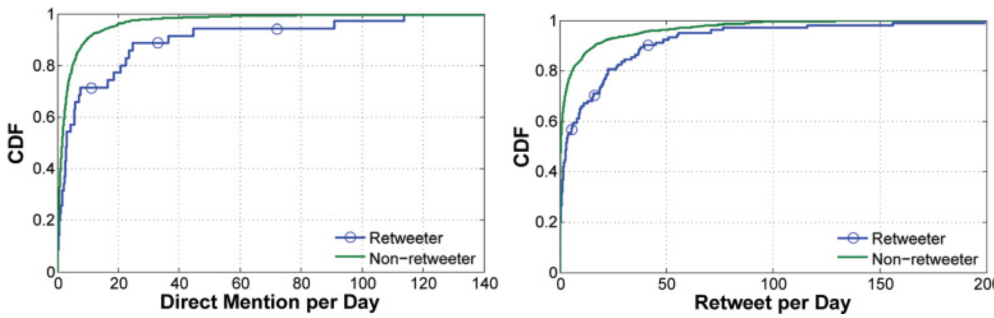

Fig. 3. CDF plots of “direct mention per day” feature in public safety training set (left figure) and “retweet per day” feature in bird flu training set (right figure).

Next, we have analyzed how feature values of the significant features are different between retweeters and nonretweeters. Because of the limited space, we only show the cumulative distribution function (CDF) plots of two features that were commonly significant across both datasets. Figure 3 shows CDFs of the |direct mention| per day feature in the public safety training set and |retweets| per day in the bird flu training set. We observe that retweeters have posted a higher number of direct mentions per day and more retweet messages per day than nonretweeters. Overall, our analysis suggests that retweeters are more advanced Twitter users, since they use advanced features more frequently (e.g., inclusion of URLs and hashtags in their tweets).

5.1. Incorporating Time Constraints

While our predictive models compute a person’s likelihood to retweet upon request, it does not predict when that person will retweet. Some situations may require important messages to be spread quickly, such as emergency alerts and SOS messages, so we also explore how to predict when a person will act on the retweeting request. To do this, we examine the person’s previous temporal behavior and use this information for prediction.

In the simplest case, our model estimates the wait time for a person to respond to a retweeting request. We further assume that retweeting events follow a Poisson process during which each retweeting occurs continuously and independently at a constant average rate. We thus use an exponential distribution model to estimate a user’s retweeting wait time with a probability. The CDF of an exponential distribution is

$$f(x;\lambda) = \begin{cases} 1 - e^{-\lambda x}, & x \geq 0, \\ 0, & x < 0. \end{cases}$$

The distribution is on the interval from zero to infinite. We measure $1/\lambda$, which is the average wait time for a user based on prior retweeting wait time. For a user’s specific retweeting wait time t, our model can predict the probability of the user’s next retweeting *P(t)* within that wait time. Figure 4 shows our model with three examples. The green line with stars indicates that a person’s average wait time is 180 minutes based on past retweeting behavior. The retweeting probability within 200 minutes is larger than 0.6. The lower a person’s average retweeting wait time t is, the higher the probability that his or her retweeting is within time t.

In practice, given a specific time constraint t, we select a *cutoff probability* c that is then used to select people whose probability of retweeting within time t is greater than or equal to c. For example, with the cutoff probability of 0.7, our model will select only those who have at least a 70% chance to retweet within the given time constraint. Incorporating the time estimation with our prediction models, we contact only people

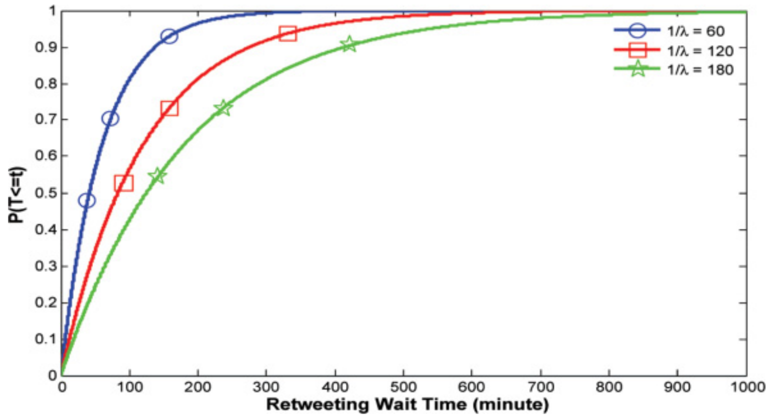

Fig. 4. Three examples of the exponential distribution.

ALGORITHM 1: Retweeter Identification Under a Time Constraint

Given a user set U, a time constraint t, cutoff probability c,
for $u \in U$ **do**
  if $classify(u)$ = "retweeter" **then**
    if $CumulativeProbOfWaitTime(u, t) >= c$ **then**
      ask u to retweet a message
    end if
  end if
end for

who are likely to retweet and whose cumulative probability of the retweeting wait time is greater than or equal to the *cutoff probability c* as described in Algorithm 1.

5.2. Incorporating Benefit and Cost

We have also explored the tradeoffs between the cost of contacting a user and the benefit of a retweet. We assume the benefit is the number of people who are directly exposed to the message as a result of the retweets, which is the total number of followers of the retweeter. Using this assumption, if our system contacts N users and K retweet, the total benefit is then the sum of all followers of the K users. Assuming a unit cost per contact, the total cost is then N. We normalize the total benefit by total cost to compute the *unit info reach per person*:

$$unit\ info\ reach\ per\ person = \frac{\sum_1^K followers(i)}{N}$$

To address the case that the same person follows multiple retweeters, we count just the number of *distinct* followers for each retweeter.

Note that measuring the effectiveness of a retweet in multiple hops (i.e., the followers of the followers to be as distinct as possible, or quantifying the influence of a node in a social network based on a recursive formulation that further investigates the influence of the followers) can be an alternative numerator of the previous evaluation metrics. But it requires collecting each follower's follower list, and even collecting additional information of followers in multiple hops. Since Twitter has changed their API limits, this approach takes a long time and may be unrealistic. Instead, we used the unit info

Table V. Prediction Accuracy (Public Safety)

Classifier	AUC	F1	F1 of Retweeter
Basic			
Random Forest	0.638	0.958	0
Naïve Bayes	0.619	0.939	0.172
Logistic	0.640	0.958	0
SMO	0.500	0.96	0
AdaBoostM1	0.548	0.962	0.1
SMOTE			
Random Forest	0.606	0.916	0.119
Naïve Bayes	0.637	0.923	0.132
Logistic	0.664	0.833	0.091
SMO	0.626	0.813	0.091
AdaBoostM1	0.633	0.933	0.129
Cost Sensitive (Weighting)			
Random Forest	**0.692**	0.954	0.125
Naïve Bayes	0.619	0.93	0.147
Logistic	0.623	0.938	0.042
SMO	0.633	0.892	0.123
AdaBoostM1	0.678	0.956	0.133

reach per person to measure how many users were directly exposed by each retweet in the first hop.

5.3. Experiments

We designed and conducted an extensive set of experiments to measure the performance of various prediction models. We also compared the effectiveness of our approach with two baselines in various conditions.

5.3.1. Evaluating Retweeter Prediction. To evaluate the performance of our prediction models, we used only the significant features found by our feature analysis (Tables III and IV) in our experiments.

Accuracy Metrics. We use three metrics to assess prediction accuracy: area under the ROC curve (AUC), F1, and F1 of the retweeter class. We use AUC as our primary performance measure, since a higher AUC means that a model is good at correctly predicting both class instances regardless of class imbalance [Fawcett 2006]. We report an overall F1 score as a reference measure, and F1 of the retweeter class on the performance of predicting minority class instances.

Settings. We ran all five prediction models under three settings: basic, SMOTE, and cost sensitive. The *basic* setting did not handle class imbalance. SMOTE was an oversampling approach in which we oversampled the minority class instances in the training set such that there was an equal number of majority and minority class instances. Under the *cost-sensitive* setting, we used a weighting scheme that weighted the minority class instances higher than the majority class instances. In our experiments, we tried five different weight ratios from 10:1 through 50:1 at intervals of 10. With five prediction models under three settings, we ran a total of 35 experiments: five in the basic setting, five in the SMOTE setting, and 25 using the cost-sensitive setting (five models by five weight ratios).

Prediction Results. Table V shows the results for the public safety dataset. Overall, the cost-sensitive setting (weighting) yielded better performance than SMOTE for both

Table VI. Prediction Accuracy (Bird Flu)

Classifier	AUC	F1	F1 of Retweeter
Basic			
Random Forest	0.707	0.877	0.066
Naïve Bayes	0.670	0.834	0.222
Logistic	0.751	0.878	0.067
SMO	0.500	0.876	0
AdaBoostM1	0.627	0.878	0.067
SMOTE			
Random Forest	0.707	0.819	0.236
Naïve Bayes	0.679	0.724	0.231
Logistic	0.76	0.733	0.258
SMO	0.729	0.712	0.278
AdaBoostM1	0.709	0.837	0.292
Cost Sensitive (Weighting, showing the best results in each model)			
Random Forest	**0.785**	0.815	0.296
Naïve Bayes	0.670	0.767	0.24
Logistic	0.735	0.742	0.243
SMO	0.676	0.738	0.256
AdaBoostM1	0.669	0.87	0.031

AUC and F1 of the retweeter class. Both Random Forest and AdaBoostM1 performed particularly well under the cost-sensitive setting. We found the similar results using the bird flu dataset (Table VI). The class imbalance problem can be observed in the poor results under the basic setting. For example, SMO completely failed to predict retweeter instances (F1 of retweeter is 0). Although both SMOTE and the cost-sensitive settings outperformed the basic one, we did not observe any clear advantage of one over the other. Note that we also ran the same experiments under a threefold cross-validation setting, and we got consistent results (i.e., SMOTE and the cost-sensitive settings outperformed the basic one).

In summary, we have found prediction configurations that produced good results by the measures of AUC and F1. Since Random Forest in the cost-sensitive setting performed the best, we used it in the rest of our experiments.

5.3.2. Comparison with Two Baselines. To validate how well our prediction approach helps improve the retweeting rate in practice, we compared the retweeting rates produced by our approach with those of two baselines: random people contact and popular people contact.

The *random people contact* approach randomly selects and asks a subset of qualified candidates on Twitter (e.g., people living in San Francisco or who tweeted about bird flu) to retweet a message. This is precisely the approach that we used during our data collection to obtain the retweeting rates for both datasets. The *popular people contact* approach first sorts candidates in our testing set by their follower count in descending order. It then selects and contacts "popular" candidates whose follower count is greater than a threshold. In our experiment, we chose 100 as the threshold since a recent study reported that more than 87% of Twitter users have fewer than 100 followers.[5] We also considered other threshold values (e.g., 50, 500, 1,000) and found that their retweeting rates were comparable.

Table VII shows the comparison of retweeting rates produced by our approach against the two baselines. Overall, our approach produced a significantly higher retweeting

[5]http://www.beevolve.com/twitter-statistics/.

Table VII. Comparison of Retweeting Rates

Approach	Retweeting Rate in Testing Set	
	Public Safety	Bird Flu
Random people contact	2.6%	8.3%
Popular people contact	3.1%	8.5%
Our prediction approach	**13.3%**	**19.7%**

Table VIII. Comparison of Retweeting Rates with Time Constraints

Approach	Average Retweeting Rate in Testing Set Under Time Constraints	
	Public Safety	Bird Flu
Random people contact	2.2%	6.5%
Popular people contact	2.7%	6.4%
Our prediction approach	**13.3%**	**13.6%**
Our prediction approach + wait-time model	**19.3%**	**14.7%**

rate than both baselines. Specifically, ours increases the average retweeting rate of two baselines by 375% (13.3% vs. 2.8%) in the public safety domain and by 135% (19.7% vs. 8.4%) in the bird flu scenario.

Adding Wait Time Constraint. We also tested our wait-time model that predicts when a person would retweet after receiving a request. We compared the retweeting rate obtained using our approach with the wait-time model with that of three settings: (1) random user contact, (2) popular user contact, and (3) our approach without the use of the wait-time model. In this experiment, the retweeting rate was the ratio of the people who retweeted our messages *within* the allotted time and the total number of people whom we contacted. In other words, if a person retweeted a requested message after the allotted time (e.g., 24 hours), he or she would be considered a nonretweeter as he or she did not meet the time constraint.

In our approach with the wait-time model, we set the cutoff probability at 0.7. As described previously, we first selected a subset of people who were predicted as retweeters and then eliminated those whose estimated probability to retweet within the given time window was smaller than the cutoff probability. We experimented with different time windows, such as 6, 12, 18, or 24 hours. Table VIII shows our experimental results with the averaged retweeting rates obtained for both of our datasets. Overall, our approach with the wait-time model outperformed the other three approaches in both datasets, achieving a 19.3% and 14.7% retweeting rate, respectively. Specifically, our model with the wait-time constraint increases the average retweeting rate of two baselines by 680% (19.3% vs. 2.45%) in the public safety domain and by 130% (14.7% vs. 6.45%) in the bird flu scenario. There is also an improvement of 45% (19.3% vs. 13.3%) in the public safety domain and 8% (14.7% vs. 13.6%) in the bird flu domain over our own algorithm when the wait-time model was not used. In summary, the *combined approach* of using our prediction model and wait-time estimation further improved retweeting rates.

Effects of Benefit and Cost. As described previously, another method of evaluating the performance of our work is via a benefit/cost analysis using the notion of information reach. We compared the results obtained during data collection with the results of our best prediction results on the testing set. Table IX shows the comparison of random user contact, popular user contact, and our approach without or with the wait-time model. The results show that our approach with/without the wait-time model achieved a higher unit info reach per person than the two baselines. In particular, our approach

Table IX. Comparison of Information Reach

Approach	Unit Info Reach Per Person	
	Public Safety	Bird flu
Random people contact	6	85
Popular people contact	11	116
Our prediction approach	**106**	**135**
Our prediction approach + wait−time model	**153**	**155**

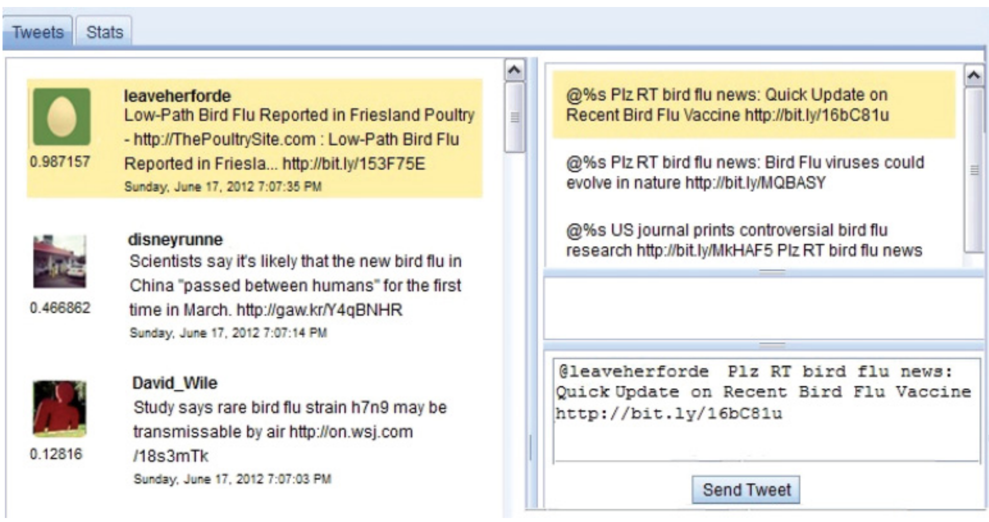

Fig. 5. The interface of our retweeter recommendation system: left panel: system-recommended candidates; right panel: a user can edit and compose a retweeting request.

with the wait-time model increased the average information reach of two baselines by 1,700% (153% vs. 8.5% = avg. (6, 11)) in the public safety domain and 54% (155% vs. 100.5% = avg. (85, 116)) in the bird flu domain, respectively.

6. REAL-TIME RETWEETER RECOMMENDATION

As mentioned earlier, our goal is to automatically identify and engage the right strangers at the right time on social media to help spread intended messages within a given time window. We thus have developed an interactive recommender system that uses our prediction model and the wait-time estimation model in *real time* to recommend the right candidates to whom retweeting requests will be sent. Figure 5 shows the interface of our system. Our system monitors the Twitter live stream and identifies a set of candidates who have posted content relevant to the topic of a retweet request (e.g., “bird flu” alerts). Such content filtering can be done by using the approaches detailed in Chen et al. [2013]. Based on the identified candidates, our system uses the prediction model to compute the candidates’ likelihood of retweeting and their probability of retweeting within the given time window t. It then recommends the top-N ranked candidates whose probability of retweeting within t is also greater than or equal to the cutoff probability c (Figure 5(a)). A user (e.g., an emergency worker) of our system can interactively examine and select the recommended candidates and control the engagement process, including editing and sending the retweeting request (Figure 5(b)).

7. MAXIMIZING RETWEETING RATE

In the previous sections, we have described classifiers we trained to model the likelihood to retweet. Such classifiers can classify a user as either a retweeter or nonretweeter. Based on the classification output, our system can contact people who are classified as retweeters. However, depending on the applications, such people selection is often associated with certain objectives, such as maximizing the retweet rate.

Since a classifier is not always perfect (it may not be 100% accurate), selecting people using classification output may not fulfill such an objective. For example, a classifier may predict user A as a retweeter with 60% retweeting probability, but she may not actually retweet a message and thus be a nonretweeter. In such scenarios, where classification results may be inaccurate, selecting people based on classification output may not fulfill our objective.

Selecting people based on their probability of retweeting instead of merely a classification result is another possibility. For example, the people selection approach could select the top-K percent users based on their probability of retweeting or it could select users for whom the probability of retweeting is above a certain threshold (e.g., more than 60%). However, in such an approach, the value of K or threshold probability has to be manually selected, which is ad hoc and not generalizable. In addition, due to the inherent imperfections (prediction errors) in any classification models, in reality, the predicted top-K people may not necessarily be the best choices in terms of maximizing the overall retweet rate. To maximize the overall retweet rate for a given set of available people, we thus have implemented the following algorithm.

In this section, we propose a *subset-selection*-based algorithm that *automatically* selects a subset of people from a set of available people using probabilities predicted by the classifier and an estimation set from which the best interval for people selection is estimated. Our algorithm is based on the assumption that the retweeting probability of a user predicted by the classifier has high correlation with the user's actual retweeting behavior. Thus, a set of users may be sorted according to their retweeting probabilities, and an optimal interval from the sorted set corresponding to the maximum retweeting rate can be estimated. Such an optimal interval is applied to any arbitrary set of available people to select a subset of them to contact. In the experimental section, we demonstrate that our algorithm is effective in maximizing the retweet rate.

Specifically, our approach for maximizing the retweet rate is shown in Figure 6: First, we divide the testing set into an estimation set (one-third of the data of the original testing set) and new testing set (two-thirds of the data of the original testing set). Next, we compute the retweeting probability of each user in the estimation set using our trained classifier. We sort the estimation set using the computed probabilities.

Let $\{\boldsymbol{p}_1, \ldots, \boldsymbol{p}_n\}$ denote the sorted list of people in the estimation set. Next, we find an interval $[i, j]$ $(1 \leq i < j \leq n)$ from this set, where the corresponding interval subset $\{\boldsymbol{p}_i, \ldots, \boldsymbol{p}_j\}$ has a maximal retweeting rate among all interval subsets. We have also tried a slight variant of this approach, which searches for a restricted choice of intervals—only those that extend to the top, that is, of the form $[i, n]$. However, that produced a suboptimal result. The best subinterval $[i_r, j_r]$ in the estimation set defines a corresponding subinterval $[i_s, j_s]$ in the new testing set, based on percentiles. That is, if m is the cardinality of the new testing set, then $i_s = [(i_r \cdot m)/n]$ and $j_s = [(j_r \cdot m)/n]$.

For example, the estimation set consists of 100 users ($n = 100$); the new testing set consists of 200 users ($m = 200$). The goal is to find the optimal interval returning the maximal retweet rate. Our algorithm measures retweets of every k% interval in the estimation set, starting from the top (n in the figure) and moving down 1% in each iteration. Say k is 10%. The first interval in the estimation set is [100, 91], the next one is [99, 90], and so on. The algorithm finds the interval returning the highest retweeting rate in the estimation set. Then, it selects users in the same interval in the new testing

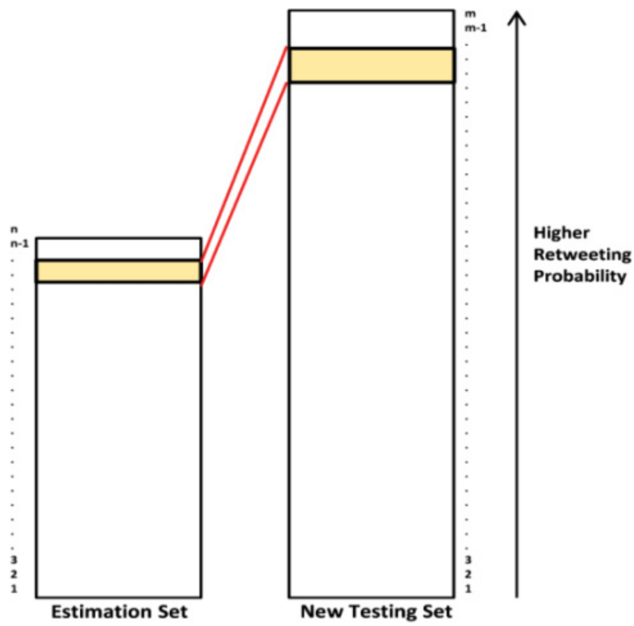

Fig. 6. Optimal interval returning maximal retweet rate in the estimation set and new testing set (colored rectangle is the optimal interval).

set. Say the optimal interval is [98, 89] in the estimation set. We select [196, 177] in the new testing set (again, $m = 200$) and contact the users in the interval. The example shows the optimal interval with 10% interval size. However, we can apply this method for other interval sizes (say, from 1% to 30%) in order to find the best interval size returning the optimal interval with the maximal retweet rate.

We can also incorporate additional constraints in our optimal interval selection. For example, we can specify the exact size of the interval or the minimum or maximum size of the interval as constraints. For example, if a minimum size of the interval is specified, our method will ignore intervals that are smaller than the specified minimum.

7.1. Experiments

We have experimentally found some reasonable classification settings for classifying a user as a retweeter or nonretweeter. As described in the previous sections, we use the probabilities computed by the trained classifier (cost sensitive with Random Forest) with our algorithm for maximizing the retweeting rate. We keep the training set (containing two-thirds of the data) and divide the original testing set into an estimation set (containing one-third of the original testing set) and new testing set (containing two-thirds of the original testing set).

To measure the performance of our algorithm, we computed the retweeting rate by increasing an interval size from 1% to 5% by 1% within the top 15% percentiles. Tables X and XI present the optimal retweeting rate of different interval sizes in the public safety and bird flu datasets. The average retweeting rate of our algorithm in the new testing set of the public safety and bird flu datasets are 16% and 54%, respectively.

Next, we compare the average retweeting rates of our classification- and subset-selection-based prediction approach with two baselines (random people contact and popular people contact) and our classification-based prediction approach. Table XII shows the comparison of retweeting rates in each of the approaches. The baselines

Table X. Variations of Optimal Interval, Estimation, and New Testing Set Retweeting Rate with Increasing Interval Size (Public Safety Dataset)

Interval Size	Optimal Interval	Retweeting Rate in Estimation Set	Retweeting Rate in New Testing Set
1%	top 0%–1%	50%	25%
2%	top 0%–2%	25%	13%
3%	top 0%–3%	17%	8%
4%	top 0%–4%	25%	19%
5%	top 0%–5%	20%	14%
Average		27%	**16%**

Table XI. Variations of Optimal Interval, Estimation, and New Testing Set Retweeting Rate with Increasing Interval Size (Bird Flu Dataset)

Interval Size	Optimal Interval	Retweeting Rate in Estimation Set	Retweeting Rate in New Testing Set
1%	top 8%–9%	100%	100%
2%	top 7%–9%	75%	56%
3%	top 6%–9%	50%	39%
4%	top 5%–9%	38%	41%
5%	top 4%–9%	30%	33%
Average		58%	**54%**

Table XII. Comparison of Retweeting Rates

	Retweeting Rate in New Testing Set	
Approach	Public Safety	Bird Flu
Random people contact	2.6%	8.4%
Popular people contact	3.1%	8.5%
Our prediction approach (classification)	11.1%	20%
Our prediction approach (classification + subset selection)	**16%**	**54%**

resulted in the lowest retweeting rate in both datasets. Our prediction approach (classification + subset selection) produced a significantly higher retweeting rate than both baselines. Specifically, ours increases the average retweeting rate of two baselines by 460% (16% vs. 2.8%) in the public safety domain and by 540% (54% vs. 8.4%) in the bird flu scenario. There is also an improvement of 45% (16% vs. 11.1%) in the public safety domain and 170% (54% vs. 20%) in the bird flu domain over our own algorithm (classification). In summary, our classification- and subset-selection-based prediction approach improved retweeting rates.

Adding Wait-Time Constraint. In order to predict when a user retweets information after he or she receives a message, we use an exponential distribution model as we described in Section 5.1.

Like our previous experiments, the retweeting rate obtained using our prediction approach (classification + subset selection) with a time constraint is compared against the baselines and our classification approach. For the baselines and our classification approach, the retweeting rate is computed as the ratio of the users who actually retweeted our messages within that time and the total number of users we actually contacted.

For our prediction approach (classification + subset selection) with the waittime model, we set the cutoff probability in the exponential distribution model to 0.7. We eliminated the users for whom the estimated wait time for the next retweeting was smaller than this cutoff probability from both of our estimation and the new testing set.

Table XIII. Comparison of Retweeting Rates Under Time Constraints

Approach	Average Retweeting Rate in New Testing Set Under Time Constraints	
	Public Safety	Bird Flu
Random people contact	2.2%	6.2%
Popular people contact	2.9%	5.6%
Our prediction approach (classification)	11.1%	12.3%
Our prediction approach (classification + subset selection) + wait-time model	**19.3%**	**19.8%**

Table XIV. Comparison of Information Reach

Approach	Unit Info Reach per User in New Testing Set	
	Public Safety	Bird flu
Random people contact	9	91
Popular people contact	14	127
Our prediction approach (classification)	152	111
Our prediction approach (classification + subset selection)	**162**	**233**

We experimented with different time windows, such as 6, 12, 18, or 24 hours. Table XIII shows our experimental results with the averaged retweeting rates obtained for both of our datasets. Overall, our prediction approach (classification + subset selection) with the wait-time model outperformed the other three approaches in both datasets, achieving a 19.3% and 19.8% retweeting rate, respectively. Specifically, our model with the wait-time constraint increases the average retweeting rate of two baselines by 640% (19.3% vs. 2.6%) in the public safety domain and by 230% (19.8% vs. 5.9%) in the bird flu scenario. There is also an improvement of 74% (19.3% vs. 11.1%) in the public safety domain and 61% (19.8% vs. 12.3%) in the bird flu domain over our own prediction approach (classification) when the wait-time model was not used. In summary, the *combined approach* of using our prediction model (classification + subset selection) and wait-time estimation further improved the retweeting rates.

Effect of Benefit and Cost. As described in Section 5.2, we have investigated the use of the benefit of retweeting and the cost of contacting a user using the notion of information reach. Thus, we compared information reach using our prediction approach (classification + subset selection), our classification approach, and the baselines. Table XIV shows the comparisons of random user contact, popular user contact, classification approach, and our prediction approach (classification + subset selection). The results show that our prediction approach (classification + subset selection) achieved a higher unit info reach per person than the two baselines. In particular, our approach with the subset selection increased the average information reach of two baselines by 1,300% (162 vs. 11.5 = avg. (9,14)) in the public safety domain and 110% (233 vs. 109 = avg. (91,127)) in bird flu domain, respectively. There is also an improvement of 7% (162 vs. 152) in the public safety domain and 110% (233 vs. 111) in the bird flu domain over our own prediction approach (classification).

To summarize, as long as we have increased the retweeting rate or found the best range giving us the highest retweeting rate, we can get a higher unit info reach per user than the baselines and the classification approach.

8. LIVE EXPERIMENTS

To validate the effectiveness of our approaches, the classification approach and classification + subset selection approach, in a *live* setting, we used our recommender

Table XV. Comparison of Retweeting Rates in Live Experiment

Approach	Retweeting Rate
Random people contact	4%
Popular people contact	9%
Our prediction approach (classification)	19%
Our prediction approach (classification + 24% subset selection)	18%
Our prediction approach (classification + 5% subset selection)	38%

Table XVI. Comparison of Retweeting Rates in Live Experiment (with Time Constraints)

Approach	Average Retweeting Rate
Random people contact	4%
Popular people contact	8.7%
Our prediction approach (classification)	18%
Our prediction approach (classification) + wait-time model	18.5%
Our prediction approach (classification + 24% subset selection) + wait-time model	21.3%
Our prediction approach (classification + 5% subset selection) + wait-time model	34.6%

system, which was developed in Sections 5, 6, and 7, to test our approaches against the two baselines (random people contact and popular people contact). First, we randomly selected 426 candidates who had recently tweeted about "bird flu" during July 2013. We then used each approach to select 100 users among the candidates. The popular people contact and our classification approach selected the top 100 candidates based on their popularity (number of followers) rank and our prediction rank, respectively. Our classification + subset selection approach also selected the best 24% range (100 candidates) among 426 candidates sorted by our prediction rank. We selected the 24% range giving us the highest retweeting rate in the estimation set as we described in Section 7. We also selected the best 5% range to see whether the best 5% range by our classification + subset selection approach would achieve a higher retweeting rate than the 24% range. If a person happened to be selected by more than one approach, we contacted the person only once to avoid overburdening the person. Overall, we contacted a total of 236 unique people. Table XV shows the comparison of retweeting rates for each approach. Our approaches outperformed two baselines in a live setting significantly. Specifically, our classification approach and classification + 24% subset selection approach increases the average retweeting rate of two baselines by more than 190% (19% vs. 6.5%) and 175% (18% vs. 6.5%), respectively. Interestingly, the classification + 5% subset selection approach even increases the average retweeting rate of two baselines by more than 480% (38% vs. 6.5%). We checked the social graph of the retweeters (those who retweeted our message). They were not connected at all. Thus, our result was unlikely to be affected by their social relationship.

We also wanted to investigate the effectiveness of our approach with time constraints. Thus, we repeated the previous experiment with different time windows, such as 6, 12, 18, or 24 hours, as we did in Sections 5.3 and 7.1. Table XVI shows the comparison of retweeting rates for each approach. Again, our approaches with our wait-time model outperformed the other three approaches. Specifically, our classification approach with the wait-time model increased the average retweeting rate of two baselines by more than 190% (18.5% vs. 6.35%). Our classification + 24% subset selection approach and classification + 5% subset selection approach with the wait-time model increased the average retweeting rate of two baselines by more than 235% (21.3% vs. 6.35%) and 440% (34.6% vs. 6.35%), respectively. Our approaches with the wait-time model outperformed our classification approach when the wait-time model was not used. In summary, this result confirms that our approaches consistently outperformed others in a live setting by a large margin.

9. DISCUSSION

Here, we discuss several of the observations during our investigation and the limitations of our current work.

9.1. Why People Retweet at a Stranger's Request

Although previous studies discuss various reasons why people retweet in general [Boyd et al. 2010; Starbird and Palen 2010], they focus on people's voluntary retweeting behavior. We were curious to find out why people retweet upon the request of a stranger. We randomly selected 50 people who retweeted per our request and asked them why they chose to retweet. Thirty-three out of 50 replied to us. Their responses revealed several reasons why people accept our retweeting requests. One reason was the trustworthiness of the content to be spread: "Because it contained a link to a significant report from a reputable media news source." Another reason is content relevance, for example, messages about their own local area: "Because it happened in my neighborhood." Interestingly, several mentioned that they retweeted because the message contained valuable information and was helpful to society: "my followers should know this or they may think this info is valuable." Some of the other reasons, such as to spread tweets to new audience or to entertain a specific audience, were discussed by others [Boyd et al. 2010] but not mentioned in our context. In the future, it would be interesting to study whether including the rationale in a retweeting request would help motivate the target strangers and affect the retweeting rate.

9.2. Retweeting with Modification

We have observed that some people retweeted our messages with modifications (e.g., adding hashtags to clarify the message or their own opinion to the original message):

> ***#publichealth*** *news: The Evolution of Bird Flu, and the Race to Keep Up http://nyti.ms/Qf6zsM @nytimesscience*
>
> ***what a shame + waste of tax $$*** *"@BayPublicSafety: @esavestheworld "Hacker created fake Sierra LaMar posting http://bit.ly/Leaojo" Plz RT"*

Such behavior suggests that the target information propagators may augment/alter the original message with additional information including their personal opinions, especially if they strongly agree/disagree with the intended information. Based on this observation, it would be interesting to investigate the additional gains and risks that a potential information propagator might bring when asked to spread the message. For example, the added hashtag (#publichealth) in the retweet listed previously would help propagate the message not only to the followers but also to those who follow the hashtag. On the opposite end, a propagator's negative opinions may affect the spread and perception of the intended message.

9.3. Generalizability

We wanted to examine how well our findings can be generalized across topics. We ran an experiment where we combined the training and testing sets of public safety and bird flu. We trained prediction models on the combined training set using the significant features identified for the combined set. The AUC in this experiment was 0.736, better than the original public safety result (0.692) but lower than the original bird flu result (0.785). The resulted retweeting rate was 12.5%, better than the random user contact (5.5%) and popular user contact (6%) for the combined set but lower than the rates achieved in public safety (13.3%) and bird flu alone (19.7%). Our results suggest that it is feasible to build a domain-independent prediction model if we have sufficient training data from different domains. We are investigating the applicability of our models to new domains, for example, new topics that our model is not trained on.

9.4. Maximizing Information Reach

In the previous section, we presented an experiment that incorporated cost/benefit and showed that our algorithm that maximizes the retweet rate achieved higher information reach than the classification approach and baselines. Another objective could be to design an algorithm that can maximize the information reach. We attempted to develop such an algorithm using a similar approach as our algorithm for maximizing the retweet rate. In particular, we tried two variants of our heuristic when finding the optimal interval from the estimation set: (1) sorting each user by the product of retweeting probability and number of followers and finding the optimal range based on the unit info reach per user and (2) sorting each user by only retweeting probability and finding the optimal range based on the unit info reach per user. However, in practice, we discovered that these heuristics produced a suboptimal result in comparison to our original heuristic (sorting by retweeting probability and finding interval that maximizes the retweet rate) for maximizing information reach. One reason these heuristics produced a suboptimal result is that both of them could be affected by users with a large number of followers. Such users may have a low retweet probability (making the first heuristic ineffective) or may not appear uniformly in both the estimation and testing set (making the second heuristic ineffective). We plan to further investigate such issues and develop solutions for maximizing information reach. Furthermore, we plan to experiment with information reach scenarios when the benefit of a retweet is computed not only from followers of those who retweeted but also from users who are indirectly exposed to the message (such as followers' followers) as a result of a retweet.

9.5. Optimizing Multiple Information Spreading Objectives

Currently, our work focuses on maximizing the retweeting rate in information diffusion. However, in practice, there may be multiple objectives to be satisfied, such as maximizing the expected net benefit or minimizing the reach time. We thus are investigating a model that can optimize multiple objectives at the same time. However, this is nontrivial, as satisfying one objective may influence the other, especially in a real-world situation, where many of these objectives may be dynamically changing (e.g., the availability of retweeting candidates and the required time frame for a message to reach a certain audience).

10. CONCLUSIONS

In this article, we have presented a feature-based prediction model that can automatically identify the right individuals at the right time on Twitter who are likely to help propagate messages per a stranger's request. We have also described a time estimation model that predicts the probability of a person to retweet the requested message within a given time window. In addition, we have developed a subset selection model to maximize the rate of retweeting and demonstrated how such a model can work under different constraints. Based on these three models, we built an interactive retweeter recommender system that allows a user to identify and engage strangers on Twitter who are most likely to help spread a message. To train and test our approaches, we collected two ground-truth datasets by *actively* engaging 3,761 people on Twitter on two topics: public safety and bird flu. Through an extensive set of experiments, we found that our approaches were able to at least *double* the retweeting rates over two baselines. With our time estimation model, our approach also outperformed other approaches significantly by achieving a much higher retweeting rate within a given time window. Furthermore, our approach has achieved a higher unit information reach per person than the baselines. In a live setting, our approach consistently outperformed the two baselines by almost doubling their retweeting rates. Overall, our approach effectively identifies qualified candidates for retweeting a message within a given time window.

ACKNOWLEDGMENTS

Research was sponsored by the U.S. Defense Advanced Research Projects Agency (DARPA) under the Social Media in Strategic Communication (SMISC) program, Agreement Number W911NF-12-C-0028. The views and conclusions contained in this document are those of the author(s) and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Defense Advanced Research Projects Agency or the U.S. government. The U.S. government is authorized to reproduce and distribute reprints for government purposes notwithstanding any copyright notation hereon.

REFERENCES

Nitin Agarwal, Huan Liu, Lei Tang, and Philip S. Yu. 2008. Identifying the influential bloggers in a community. In *WSDM*.

Eytan Bakshy, Jake M. Hofman, Winter A. Mason, and Duncan J. Watts. 2011. Everyone's an influencer: Quantifying influence on Twitter. In *WSDM*.

Eytan Bakshy, Itamar Rosenn, Cameron Marlow, and Lada Adamic. 2012. The role of social networks in information diffusion. In *WWW*.

Danah Boyd, Scott Golder, and Gilad Lotan. 2010. Tweet, tweet, retweet: Conversational aspects of retweeting on twitter. In *HICSS*.

Ceren Budak, Divyakant Agrawal, and Amr El Abbadi. 2011. Limiting the spread of misinformation in social networks. In *WWW*.

Meeyoung Cha, Hamed Haddadi, Fabricio Benevenuto, and Krishna P. Gummadi. 2010. Measuring user influence in twitter: The million follower fallacy. In *ICWSM*.

Vineet Chaoji, Sayan Ranu, Rajeev Rastogi, and Rushi Bhatt. 2012. Recommendations to boost content spread in social networks. In *WWW*.

Nitesh V. Chawla, Kevin W. Bowyer, Lawrence O. Hall, and W. Philip Kegelmeyer. 2002. SMOTE: Synthetic minority over-sampling technique. *J. Artif. Int. Res.* 16, 1 (June 2002), 321–357.

Jilin Chen, Allen Cypher, Clemens Drews, and Jeffrey Nichols. 2013. CrowdE: Filtering tweets for direct customer engagements. In *ICWSM*.

Kailong Chen, Tianqi Chen, Guoqing Zheng, Ou Jin, Enpeng Yao, and Yong Yu. 2012. Collaborative personalized tweet recommendation. In *SIGIR*.

Paul T. Costa and Robert R. McCrae. 1992. *Revised NEO Personality Inventory (NEO PI-R) and NEP Five-factor Inventory (NEO-FFI): Professional Manual.* Psychological Assessment Resources.

Lisa A. Fast and David C. Funder. 2008. Personality as manifest in word use: Correlations with self-report, acquaintance report, and behavior. *J. Personality Social Psychol.* 94, 2 (2008), 334.

Tom Fawcett. 2006. An introduction to ROC analysis. *Pattern Recogn. Lett.* 27, 8 (June 2006), 861–874.

Wei Feng and Jianyong Wang. 2013. Retweet or not? Personalized tweet re-ranking. In *WSDM*.

Alastair J. Gill, Scott Nowson, and Jon Oberlander. 2009. What are they blogging about? Personality, topic and motivation in blogs. In *ICWSM*.

Amit Goyal, Francesco Bonchi, and Laks V. S. Lakshmanan. 2010. Learning influence probabilities in social networks. In *WSDM*.

Mark Hall, Eibe Frank, Geoffrey Holmes, Bernhard Pfahringer, Peter Reutemann, and Ian H. Witten. 2009. The WEKA data mining software: An update. *SIGKDD Explor. Newsl.* 11, 1 (Nov. 2009), 10–18.

Tuan-Anh Hoang and Ee-Peng Lim. 2012. Virality and susceptibility in information diffusions. In *ICWSM*.

Nathan O. Hodas and Kristina Lerman. 2014. The simple rules of social contagion. *Sci. Rep.* 4 (2014). http://www.nature.com/srep/2014/140311/srep04343/full/srep04343.html.

Junming Huang, Xue-Qi Cheng, Hua-Wei Shen, Tao Zhou, and Xiaolong Jin. 2012. Exploring social influence via posterior effect of word-of-mouth recommendations. In *WSDM*.

Kyumin Lee, James Caverlee, and Steve Webb. 2010. Uncovering social spammers: Social honeypots + machine learning. In *SIGIR*.

Kyumin Lee, Brian David Eoff, and James Caverlee. 2011. Seven months with the devils: A long-term study of content polluters on Twitter. In *ICWSM*.

Kyumin Lee, Jalal Mahmud, Jilin Chen, Michelle Zhou, and Jeffrey Nichols. 2014. Who will retweet this? Automatically identifying and engaging strangers on twitter to spread information. In *IUI*.

Xu-Ying Liu and Zhi-Hua Zhou. 2006. The influence of class imbalance on cost-sensitive learning: An empirical study. In *ICDM*.

Sofus A. Macskassy and Matthew Michelson. 2011. Why do people retweet? Anti-homophily wins the day! In *ICWSM*.

Jalal Mahmud, Michelle X. Zhou, Nimrod Megiddo, Jeffrey Nichols, and Clemens Drews. 2013. Recommending targeted strangers from whom to solicit information on social media. In *IUI*.

Jeffrey Nichols and Jeon-Hyung Kang. 2012. Asking questions of targeted strangers on social networks. In *CSCW*.

James W. Pennebaker, Martha E. Francis, and Roger J. Booth. 2001. *Linguistic Inquiry and Word Count*. Lawrence Erlbaum Associates.

Daniel M. Romero, Wojciech Galuba, Sitaram Asur, and Bernardo A. Huberman. 2011. Influence and passivity in social media. In *ECML/PKDD*.

Claude Shannon. 1948. A mathematical theory of communication. *Bell Syst. Tech. J.* 27 (July, October 1948), 379–423, 623–656.

Yaron Singer. 2012. How to win friends and influence people, truthfully: Influence maximization mechanisms for social networks. In *WSDM*.

Kate Starbird and Leysia Palen. 2010. *Pass It On? Retweeting in Mass Emergency.* International Community on Information Systems for Crisis Response and Management.

Greg Ver Steeg and Aram Galstyan. 2012. Information transfer in social media. In *WWW*.

Jianshu Weng, Ee-Peng Lim, Jing Jiang, and Qi He. 2010. TwitterRank: Finding topic-sensitive influential twitterers. In *WSDM*.

Yiming Yang and Jan O. Pedersen. 1997. A comparative study on feature selection in text categorization. In *ICML*.

Tal Yarkoni. 2010. Personality in 100,000 words: A large-scale analysis of personality and word use among bloggers. *J. Res. Personality* 44, 3 (2010), 363–373.

Received March 2014; revised August 2014; accepted September 2014