

# Detecting and Tracking the Spread of Astroturf Memes in Microblog Streams

Jacob Ratkiewicz\*

Michael Conover

Mark Meiss

Bruno Gonçalves

Snehal Patil

Alessandro Flammini

Filippo Menczer

Center for Complex Networks and Systems Research  
Pervasive Technology Institute  
School of Informatics and Computing, Indiana University, Bloomington, IN, USA

## ABSTRACT

Online social media are complementing and in some cases replacing person-to-person social interaction and redefining the diffusion of information. In particular, microblogs have become crucial grounds on which public relations, marketing, and political battles are fought. We introduce an extensible framework that will enable the **real-time analysis of meme diffusion in social media** by mining, visualizing, mapping, classifying, and modeling massive streams of public microblogging events. We describe a Web service that leverages this framework to track political memes in Twitter and help detect astroturfing, smear campaigns, and other misinformation in the context of U.S. political elections. We present some **cases of abusive behaviors** uncovered by our service. Finally, we discuss promising preliminary results on the detection of suspicious memes via supervised learning based on features extracted from the topology of the diffusion networks, sentiment analysis, and crowd-sourced annotations.

## Categories and Subject Descriptors

H.2.8 [Database Management]: Database applications—*Data mining*; H.3.4 [Information Storage and Retrieval]: Systems and Software—*Information networks*; H.4 [Information Systems Applications]: Miscellaneous; K.4.1 [Computers and Society]: Public Policy Issues

## General Terms

Measurement

## Keywords

Twitter, Truthy, Memes, Information Diffusion, Social Media, Microblogs, Classification, Politics

\*Corresponding author. Email: jpr@cs.indiana.edu

## 1. INTRODUCTION

Social networking and microblogging services reach hundreds of million users and have become fertile ground for a variety of research efforts, since they offer an opportunity to study patterns of social interaction among far larger populations than ever before. In particular, Twitter has recently generated much attention in the research community due to its peculiar features, enormous popularity, and open policy on data sharing.

Along with the growth in reach of microblogs, we are also observing the emergence of useful information that can be mined from their data streams [4, 48, 11]. However, as microblogs become valuable media to spread information, e.g., for marketers and politicians, it is natural that people find ways to abuse them. As a result, we observe various types of illegitimate use, such as spam [52, 21, 49]. In this paper we focus on one particular type of abuse, namely *political astroturf* — campaigns disguised as spontaneous, popular “grassroots” behavior that are in reality carried out by a single person or organization. This is related to spam but with a more specific domain context, and with potentially larger consequences. The importance of political astroturf stems from the unprecedented opportunities created by social media for increased participation and information awareness among the Internet-connected public [1, 8, 2].

Online social media tools have played a crucial role in the successes and failures of numerous political campaigns and causes, from the grassroots organizing power of Barack Obama’s 2008 presidential campaign, to Howard Dean’s failed 2004 presidential bid and the first-ever Tea Party rally [46, 39, 51]. Moreover, traditional media pay close attention to the ebb and flow of communication on social media platforms, and with this scrutiny comes the potential for these discussions to reach a far larger audience than simply the social media users.

While some news coverage of social media may seem banal and superficial, their focus is not without merit. Social media, such as Twitter, often enjoy substantial user bases with participants drawn from diverse geographic, social and political backgrounds [29, 3]. Moreover, the user-as-information-producer model provides researchers and news organizations alike a means of instrumenting and observing, in real-time, a large sample of the nation’s political participants. So relevant is this discursive space, in fact, that the Library of Congress has recently undertaken the project of archiving a complete record of the discourse produced by Twitter users [41]. Despite the benefits associated with increased information availability and grassroots political organization, the same structural and systematic properties that enable Twitter users

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$10.00.

to quickly sound the alarm about a developing emergency [16] can also be leveraged to spread lies and misinformation.

Unlike traditional news sources, social media provide little in the way of individual accountability or fact-checking mechanisms, meaning that catchiness and repeatability, rather than truthfulness, can function as the primary drivers of information diffusion in these information networks. While flame wars and hyperbole are hardly new phenomena online, Twitter’s 140-character sound-bytes are ready-made headline fodder for the 24-hour news cycle. More than just the calculated emissions of high-profile users like Sarah Palin and Barack Obama, consider the fact that several major news organizations picked up on the messaging frame of a viral tweet relating to the allocation of stimulus funds, succinctly describing a study of decision making in drug-addicted macaques as “Stimulus \$ for coke monkeys” [47].

While the “coke monkeys” meme developed organically from the attention dynamics of thousands of users, it illustrates the powerful and potentially detrimental role that social media can play in shaping the public discourse. As we will demonstrate, a motivated attacker can easily orchestrate a distributed effort to mimic or initiate this kind of organic spreading behavior, and with the right choice of inflammatory wording, influence a public well beyond the confines of his or her own social network.

Here we describe a system to analyze the diffusion of information in social media, and in particular to automatically identify and track orchestrated, deceptive efforts to mimic the organic spread of information through the Twitter network. The main contributions of this paper are:

- The introduction of an extensible framework for the real-time analysis of *meme* diffusion in social media by mining, visualizing, mapping, classifying, and modeling massive streams of public microblogging events. (§ 3)
- The design and implementation of a Web service that leverages our framework to track political memes in Twitter and help detect astroturfing, smear campaigns, and other misinformation in the context of U.S. political elections. (§ 4)
- A description and analysis of several cases of abusive behavior uncovered by our service. (§ 5)
- Promising preliminary results on the detection of suspicious memes via supervised learning, which achieve around 90% accuracy based on features extracted from the topology of the diffusion networks, sentiment analysis, and crowdsourced annotations. (§ 6)

## 2. RELATED WORK AND BACKGROUND

### 2.1 Information Diffusion

The study of opinion dynamics and information diffusion in social networks has a long history in the social, physical, and computational sciences [38, 17, 5, 13, 6, 31, 32]. While usually referred to as ‘viral’ [38], the way in which information or rumors diffuse in a network has several important differences with respect to infectious diseases. Rumors gradually acquire more credibility and appeal as more and more network neighbors acquire them. After some time, a threshold is crossed and the rumor becomes so widespread that it is considered as ‘common knowledge’ within a community and hence, true. In the case of information propagation in the real world as well as in the blogosphere, the problem is significantly complicated by the fact that the social network structure is unknown. Without explicit linkage data investigators must rely on heuristics

at the node level to infer the underlying network structure. Gomez *et al.* propose an algorithm that can efficiently approximate linkage information based on the times at which specific URLs appear in a network of news sites [19]. However, even in the case of the Twitter social network, where explicit follower/followee social relations exist, they are not all equally important [26]. Fortunately for our purposes, Twitter provides an explicit way to mark the diffusion of information in the form of *retweets*. This metadata tells us which links in the social network have actually played a role in the diffusion of information.

Additionally, conversational aspects of social interaction in Twitter have recently been studied [12, 24, 20]. For example, Mendoza *et al.* examined the reliability of retweeted information in the hours following the 2010 Chilean earthquake [35]. They found that false information is more likely to be questioned by users than reliable accounts of the event. Their work is distinct from our own in that it does not investigate the dynamics of misinformation propagation. Finally, recent modeling work taking into account user behavior, user-user influence and resource virulence has been used to predict the spread of URLs through the Twitter social network [18].

### 2.2 Mining Microblog Data

Several studies have demonstrated that information shared on Twitter has some intrinsic value, facilitating, e.g., predictions of box office success [4] and the results of political elections [48]. Content has been further analyzed to study consumer reactions to specific brands [28], the use of tags to alter content [25], its relation to headline news [30], and on the factors that influence the probability of a meme to be retweeted [45]. Other authors have focused on how passive and active users influence the spreading paths [42].

Recent work has leveraged the collective behavior of Twitter users to gain insight into a number of diverse phenomena. Analysis of tweet content has shown that some correlation exists between the global mood of its users and important worldwide events [10], including stock market fluctuations [11]. Similar techniques have been applied to infer relationships between media events such as presidential debates and affective responses among social media users [15]. Sankaranarayanan *et al.* developed an automated breaking news detection system based on the linking behavior of Twitter users [44], while Heer and Boyd describe a system for visualizing and exploring the relationships between users in large-scale social media systems [23]. Driven by practical concerns, others have successfully approximated the epicenter of earthquakes in Japan by treating Twitter users as a geographically-distributed sensor network [43].

### 2.3 Political Astroturf and Truthiness

In the remainder of this paper we describe a system designed to detect astroturfing campaigns on Twitter. As an example of such a campaign, we turn to an illustrative case study documented by Metaxas and Mustafaraj, who describe a concerted, deceitful attempt to cause a specific URL to rise to prominence on Twitter through the use of a distributed network of nine fake user accounts [36]. In total these accounts produced 929 tweets over the course of 138 minutes, all of which included a link to a website smearing one of the candidates in the 2009 Massachusetts special election. The tweets injecting this meme mentioned users who had previously expressed interest in the Massachusetts special election, being prime candidates to act as rebroadcasters. By this the initiators sought not just to expose a finite audience to a specific URL, but to trigger an information cascade that would lend a sense of credibility and grassroots enthusiasm to a specific political message. Within hours, a substantial portion of the targeted users re-

tweeted the link, resulting in rapid spreading that was detected by Google’s real-time search engine. This caused the URL in question to be promoted to the top of the Google results page for the query ‘martha coakley’ — a so-called *Twitter bomb*. This case study demonstrates the ease with which a focused effort can initiate the viral spread of information on Twitter, and the serious consequences this can have.

Our work is related to the detection of spam in Twitter, which has been the subject of several recent studies. Grier *et al.* provide a general overview of spam on Twitter [21], focusing on spam designed to cause users to click a specific URL. Grouping together tweets about the same URL into spam ‘campaigns,’ they find a minimal amount of collusion between spammer accounts. Boyd *et al.* also analyze Twitter spam with respect to a particular meme [52]. Using a hand-classified set of 300 tweets, they identify several differences between spam and good user accounts, including the frequency of tweets, age of accounts, and their respective periphery in the social graph. Benevenuto *et al.* [7] use content and user behavior attributes to train a machine learning apparatus to detect spam accounts. They build a classifier that achieves approximately 87% accuracy in identifying spam tweets, and similar accuracy in detecting the spam accounts themselves.

The mass creation of accounts, the impersonation of users, and the posting of deceptive content are all behaviors that are likely common to both spam and political astroturfing. However, political astroturf is not exactly the same as spam. While the primary objective of a spammer is often to persuade users to click a link, someone interested in promoting an astroturf message wants to establish a false sense of group consensus about a particular idea. Related to this process is the fact that users are more likely to believe a message that they perceive as coming from several independent sources, or from an acquaintance [27]. Spam detection systems often focus on the content of a potential spam message — for instance, to see if the message contains a certain link or set of tags. In detecting political astroturf, we focus on how the message is delivered rather than on its content. Further, many of the users involved in propagating a successfully astroturfed message may in fact be legitimate users, who are unwittingly complicit in the deception, having been themselves deceived by the original core of automated accounts. Thus, existing methods for detecting spam that focus on properties of user accounts, such as the number of URLs in tweets originating from that account or the interval between successive tweets, would be unsuccessful in finding such astroturfed memes.

In light of these characteristics of political astroturf, we need a definition that allows us to discriminate such falsely-propagated information from organically propagated information that originates at the real grassroots. We thus decided to borrow a term, *truthy*, to describe political astroturf memes. The term was coined by comedian Stephen Colbert to describe something that a person claims to know based on emotion rather than evidence or facts. We can then define our task as the detection of truthy memes in the Twitter stream. Not every truthy meme will result in a viral information cascade like the one documented by Metaxas and Mustafaraj, but we wish to test the hypothesis that the initial stages exhibit common signatures that can help us identify this behavior.

### 3. ANALYTICAL FRAMEWORK

Social media analysis presents major challenges in the area of data management, particularly when it comes to interoperability, curation, and consistency of process. Due to diversity among site designs, data models, and APIs, any analytical tools written by researchers to address one site are not easily portable to another. To focus on the common features of all social media and microblog-

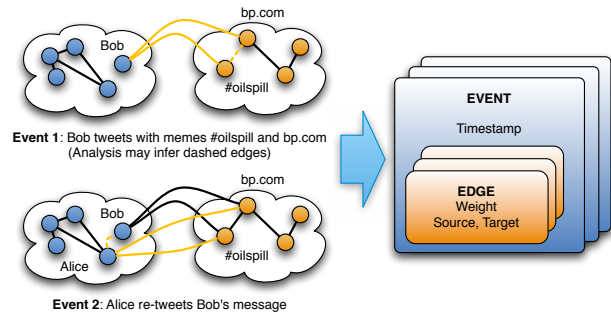


Figure 1: The Klatsch model of streaming social media events.

ging sites, we developed a unified framework, which we call *Klatsch*, that makes it possible to analyze the behavior of users and diffusion of ideas in a broad variety of data feeds. The Klatsch framework is designed to provide data interoperability for the real-time analysis of massive social media data streams (millions of posts per day) from sites with diverse structures and interfaces. To this end, we model a generic stream of social networking data as a series of events that represent interactions between actors and memes, as shown in Figure 1.


In the Klatsch model, social networking sites are sources of a timestamped series of events. Each event involves some number of actors (entities that represent individual users), some number of memes (entities that represent units of information at the desired level of detail), and interactions among those actors and memes. For example, a single Twitter post might constitute an event involving three or more actors: the poster, the user she is retweeting, and the people she is addressing. The post might also involve a set of memes consisting of ‘hashtags’ and URLs referenced in the tweet. Each event can be thought of as contributing a unit of weight to edges in a network structure, where nodes are associated with either actors or memes. This is not a strictly bipartite network: actors can be linked through replying or mentioning, and memes by concurrent discussion or semantic similarity. The timestamps associated with the events allow us to observe the changing structure of this network over time.

#### 3.1 Meme Types

To study the diffusion of information on Twitter it is necessary to single out features that can be used to identify a specific topic as it propagates through the social substrate. While there exist many sophisticated statistical techniques for modeling the underlying topics present in bodies of text [9, 14], the small size of each tweet and the contextual drift present in streaming data create significant complications [50]. Fortunately, several conventions shared among Twitter users allow us to avoid these issues entirely. We focus on the following features to identify different types of memes:

**Hashtags** The Twitter community uses tokens prefixed by a hash (#) to label the topical content of tweets. Some examples of popular tags are #gop, #obama, and #desen, marking discussion about the Republican party, President Obama, and the Delaware race for U.S. Senate, respectively. These are often called *hashtags* or *#tags*.

**Mentions** A Twitter user can call another user’s attention to a particular post by including that user’s screen name in the post, prepended by the @ symbol. These mentions can be used as a way to carry on conversations between users (*replies*),



**Figure 2: Example of a meme diffusion network involving three users mentioning and retweeting each other. The values of various node statistics are shown next to each node. The strength  $s$  refers to weighted degree.**

or to denote that a particular Twitter user is being discussed (*mentions*).

**URLs** We extract URLs from tweets by extracting strings of valid URL characters that begin with `http://`. Honeycutt *et al.* suggest that URLs are most clearly associated with the transmission of information on Twitter [24].

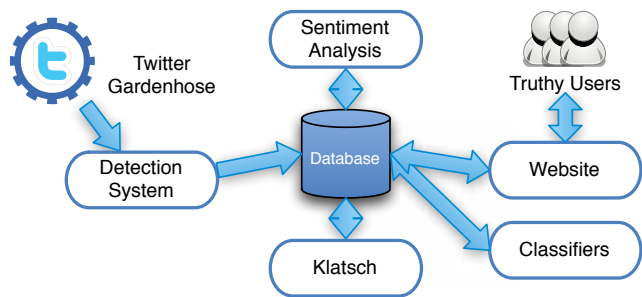
**Phrases** Finally, we consider the entire text of the tweet itself to be a meme, once all Twitter metadata, punctuation, and URLs have been removed.

Relying on these conventions we are able to focus on the ways in which a large number of memes propagate through the Twitter social network.

### 3.2 Network Edges

To represent the flow of information through the Twitter community we construct a directed graph in which nodes are individual user accounts. An example diffusion network involving three users is shown in Figure 2. An edge is drawn from node  $A$  to  $B$  when either  $B$  is observed to retweet a message from  $A$ , or  $A$  mentions  $B$  in a tweet. The weight of an edge is incremented each time we observe an event connecting two users. In this way, either type of edge can be understood to represent a flow of information from  $A$  to  $B$ . Observing a retweet at node  $B$  provides implicit confirmation that information from  $A$  appeared in  $B$ 's Twitter feed, while a mention of  $B$  originating at node  $A$  explicitly confirms that  $A$ 's message appeared in  $B$ 's Twitter feed. This may or may not be noticed by  $B$ , therefore mention edges are less reliable indicators of information flow compared to retweet edges.

We determine who was replied to or retweeted not by parsing the text of the tweet, which can be ambiguous (as in the case when a tweet is marked as being a 'retweet' of multiple people). Rather, we rely on Twitter metadata that we download along with the text of the tweet, and which designates users as being the users replied to or retweeted by each message. Thus, while the text of a tweet may contain several mentions, we only draw an edge to the user who is explicitly designated as the mentioned user by the tweet metadata. Note that this is separate from our use of mentions as memes (§ 3.1), which we parse from the text of the tweet.



**Figure 3: The Truthy system architecture.**

## 4. TRUTHY SYSTEM ARCHITECTURE

We built a system called Truthy (`truthy.indiana.edu`). A general overview of the components of our system is shown in Figure 3. Truthy includes several components: a low-level system overseeing the collecting and processing of the raw data feeds from the Twitter API, the meme detection framework, The Klatsch framework responsible for computing key network statistics and layouts, and a Web based presentation framework that allows us to collect user input on which memes the community deems most suspicious. These components are described next. Network statistics and community-generated annotations are the primary inputs to the classification apparatus discussed in § 6.


### 4.1 Streaming Data Collection

To collect meme diffusion data we rely on whitelisted access to the Twitter 'gardenhose.' The gardenhose provides detailed data on a sample of the Twitter corpus at a rate that varied between roughly 4 million tweets a day near the beginning of our study, to around 8 million tweets per day at the time of this writing. We distinguish here between the gardenhose and the firehose, the latter of which provides an unfiltered dump of all Twitter's traffic, but is only available to users who purchase access. While the process of sampling edges (tweets between users) from a network to investigate structural properties has been shown to produce suboptimal approximations of true network characteristics [33], we find that the analyses described below are able to produce accurate classifications of truthy memes even in light of this shortcoming. All collected tweets are stored in files at a daily time resolution. We maintain files both in a verbose JSON format containing all the features provided by Twitter, and in a more compact format that contains only the features used in our analysis. This collection is accomplished by a component of our system that operates asynchronously from the others.

### 4.2 Meme Detection

A second component of our system is devoted to scanning the collected tweets in real time, by pulling data from the daily files described above. The task of this meme detection component (Figure 4) is to determine which of the collected tweets are to be stored in our database and subjected to further analysis. Our goal is to collect only tweets (a) with content related to the political elections, and (b) of sufficiently general interest. We implemented a filtering step for each of these criteria, described below.

To identify politically relevant tweets, we turn to a hand-curated collection of approximately 2500 keywords relating to the 2010 U.S. midterm elections. This keyword list contains the names of all candidates running for U.S. federal office, as well as any common variations and known Twitter account usernames. The collec-




**Figure 4:** Our meme detection and tracking system consists of three separate, asynchronous components — the tweet collection, which downloads tweets and saves them to disk; the tweet filter, which determines tweets likely to relate to politics; and the meme filter, which identifies memes of significant general interest and saves them in the database.

tion further contains the top 100 hashtags that co-occurred with the hashtags #tcot and #p2 (the top conservative and liberal tags, respectively) during the last ten days of August 2010. The motivation for including explicit hashtags in the filter is not to ensure that these terms are tracked by the system (though this is a side effect), but rather to capitalize on the common behavior of Twitter users whereby they include chains of tags to identify multiple relevant topics of interest. This component, too, operates asynchronously. It is capable of processing tweets at a rate of about 10 times faster than our sampling rate, allowing it to easily handle bursts of traffic. We refer to this component as the `tweet_filter`.

Simply including all the tweets at this step would have resulted in a proliferation of distinct memes, as it would have included as a meme any hashtag, URL, username, or phrase mentioned by any user even one time. We thus implemented a second stage of filtering designed to identify those tweets containing memes of sufficiently general interest. We refer to this stage of filtering as the `meme_filter`.

The `meme_filter`, like the `tweet_filter`, reads tweets in real time. However, since tweets in the gardenhose are not guaranteed to be in strict temporal order, the `meme_filter` inserts all tweets read into a priority queue that orders them by their timestamp. Tweets are then processed in the order that they are removed from the queue. This does not guarantee that tweets will be read in sorted order, but greatly decreases the number of out-of-order tweets — for a priority queue of size  $n$ , any tweet less than  $n$  places out of order will be correctly ordered. We found empirically that  $n = 1000$  decreased out-of-order tweets to manageable levels. It is necessary to present tweets in-order to subsequent layers, to make maintenance of the sliding activation window (described next) more efficient. Thus any out-of-order tweets remaining after this step are discarded.

The `meme_filter`'s goal is to extract only those tweets that pertain to memes of significant general interest. To this end, we extract all memes (of the types described in § 3.1) from each incoming tweet, and track the activation over the past hour of each meme, in real time. If any meme exceeds a rate threshold of five mentions in a given hour it is considered 'activated;' any tweets containing that meme are then stored. If a tweet contains a meme that is already considered activated due to its presence in previous tweets, it is stored immediately. When the mention rate of the meme drops below the activation limit, it is no longer considered



**Figure 5:** The Klatsch framework architecture.

activated and tweets containing the meme are no longer automatically stored. Note that a tweet can contain more than one meme, and thus the activation of multiple memes can be triggered by the arrival of a single tweet. We chose a low rate threshold with the understanding that if a meme is observed five times in our sample it is likely mentioned many more times in Twitter at large.

The tracking of a new tweet consists of three steps: (i) removing tweets outside the current sliding activation window; (ii) extracting memes from the tweet and tracking their activation; and (iii) storing tweets related to any now activated memes. Because the tweets are presented in sorted order, and the number of memes in a tweet is bounded by the constant tweet length, step (i) can be completed in time linear in the number of old tweets, and steps (ii) and (iii) require constant time.

Prior to settling on this detection strategy for topics of general interest, we experimented with a more complicated strategy based on examining the logarithmic derivative of the number of mentions of a particular meme, computed hourly. This approach was inspired by previous work on attention dynamics in Wikipedia [40]. Since many memes with bursty behavior have low volume, we augmented the burst detection algorithm with a second predicate that included memes that appeared in a minimum percentage of the tweets over the past hour. We eventually discarded this hybrid detection mechanism due to the complexity of choosing appropriate parameters, in favor of the simpler scheme described above.

Our system has tracked a total of approximately 305 million tweets collected from September 14 until October 27, 2010. Of these, 1.2 million contain one or more of our political keywords; detection of interesting memes further reduced this set to 600,000 tweets actually entered in our database for analysis.

### 4.3 Network Analysis

The Klatsch framework is responsible for network analysis and layout for visualization of the diffusion patterns on the Truthy Web site. It consists of several components, as depicted in Figure 5. The key components of the system are: a set of input adapters for importing external social network data into the Klatsch data model; support for a variety of standard graph layout and visualization algorithms; a flexible scripting language for coding site-agnostic analysis modules; and a set of export modules, including an embedded light-weight Web server, for visualizing analysis, saving statistical results, supporting interactive Web tools, and producing publication-ready graphs and tables.

Klatsch includes an embedded domain-specific scripting language with advanced features such as first-order functions, streams, and `map/filter/reduce` primitives. For instance, the inclusion of streams as a first-order data type supports the lazy evaluation of algorithms that operate on the nodes and edges of large graphs. Our graph analysis and visualization algorithms are implemented in the Klatsch language. As an example of its expressiveness, consider the fol-

lowing problem: given a user and a meme, find out the average proportion of tweets about that meme that the user accounts for, among all tweets received by people he tweets to. (In other words, if Fred tweets a meme to Barney 3 times and Barney receives 6 tweets about that meme in total, that's 0.5.) This complex calculation can be performed by the following code snippet:

```
find_meme_prop = proc (actor_id)
  g.eo(anode(actor_id)).map(proc (e) e.w()
    / g.si(e.dst())).list().mean();
```

To characterize the structure of the diffusion network we compute several statistics based on the topology of the largest connected component of the retweet / mention graph. These include the number of nodes and edges in the graph, the mean degree and strength of nodes in the graph, mean edge weight, clustering coefficient of the largest connected components and the standard deviation and skew of each network's in-degree, out-degree and strength distributions (cf. Figure 2). Additionally we track the out-degree and out-strength of the most prolific broadcaster, as well as the in-degree and in-strength of the most focused-upon user. We also monitor the number of unique injection points of the meme in the largest connected component, reasoning that organic memes (such as those relating to news events) will be associated with larger number of originating users.

#### 4.4 Sentiment Analysis

In addition to the graph-based statistics described above we utilize a modified version of the Google-based Profile of Mood States (GPOMS) sentiment analysis method introduced by Pepe and Bollen [37] in the analysis of meme-specific sentiment on Twitter. The GPOMS tool assigns to a body of text a six-dimensional vector with bases corresponding to different mood attributes, namely *Calm*, *Alert*, *Sure*, *Vital*, *Kind*, and *Happy*. To produce scores for a meme along each of the six dimensions, GPOMS relies on a vocabulary taken from an established psychometric evaluation instrument known as POMS [34]. The original POMS test asks evaluatees to rate their emotional state with respect to a vocabulary of 72 adjectives, and associates each adjective with corresponding mood dimensions. The strength of an evaluatees identification with various adjectives contributes to summed scores along the mood dimensions.

Users of social media, however, are notoriously resourceful in their ability to create new lexicons for expression. To address this issue, the GPOMS tool relies on the Google n-gram corpus,<sup>1</sup> which identifies the co-occurrence rates of approximately 2.5 billion token pairings observed in a corpus of approximately one trillion Web pages. By associating the original POMS terms with their Google 5-gram partners of the form "I feel X and Y" where X is a POMS term, the GPOMS tool is able to expand its target lexicon to 964 tokens, each of which can be transitively associated with an underlying mood dimension. Consequently, observations of these tokens in a body of text contribute to the magnitude of a given mood dimension proportionate to the co-occurrence rate of the term and each of the original 72 POMS adjectives. The resulting mood vector is normalized to unit length, resulting in magnitudes ranging between -1 and 1 along each of the six dimensions. We applied the GPOMS methodology to the collection of tweets, obtaining a six-dimensional mood vector for each meme.

<sup>1</sup>Available at [googleresearch.blogspot.com/2006/08/all-our-n-gram-are-belong-to-you.html](http://googleresearch.blogspot.com/2006/08/all-our-n-gram-are-belong-to-you.html)

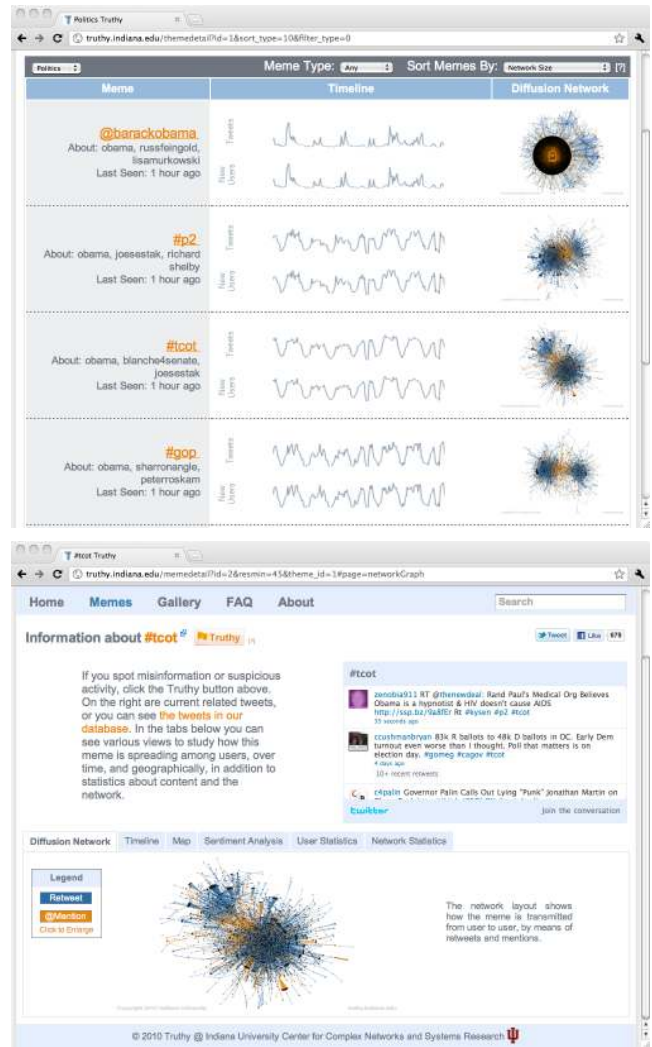


Figure 6: Screenshots of the Truthy Web site meme overview page (top) and meme detail page (bottom).

#### 4.5 Web Interface

The final component of our analytical framework includes a dynamic Web interface to allow users to inspect memes through various views, and annotate those they consider to be truthful. Raw counts of these user annotations are used as input to the classification apparatus described in § 6. To facilitate the decision making process, we provide a mixed presentation of statistical information and interactive visualizations elements. Figure 6 provide snapshots of summary and detailed views available on the Truthy site.

Users who wish to explore the Truthy database using the Web interface can sort memes according to a variety of ranking criteria, including the size of the largest connected component, number of user annotations, number of users, number of tweets, number of tweets per user, number of retweets, and number of meme injection points. This list-based presentation of memes functions as a concise, high-level view of the data, allowing users to examine related keywords, time of most recent activity, tweet volume sparklines and thumbnails of the information diffusion network. At this high level users can examine a large number of memes quickly and subsequently drill down into those that exhibit interesting behavior.

Once a user has selected an individual meme for exploration, she is presented with a more detailed presentation of statistical data and interactive visualizations. Here the user can examine the statistical data described above, tweets relating the meme of interest, and sentiment analysis data. Additionally users can explore the temporal data through an interactive annotated timeline, inspect a force-directed layout of the meme diffusion network, and view a map of the tweet geo-locations. Upon examining these features, the user is then able to make a decision as to whether this meme is truthful or not, and can indicate her conclusion by clicking a button at the top of the page.

## 5. EXAMPLES OF TRUTHY MEMES

The Truthy site allowed us to identify several truthful memes. Some of these cases caught the attention of the popular press due to the sensitivity of the topic in the run up to the political elections, and subsequently many of the accounts involved were suspended by Twitter. Below we illustrate a few representative examples.

**#ampat** The #ampat hashtag is used by many conservative users on Twitter. What makes this meme suspicious is that the bursts of activity are driven by two accounts, @CSteven and @CStevenTucker, which are controlled by the same user, in an apparent effort to give the impression that more people are tweeting about the same topics. This user posts the same tweets using the two accounts and has generated a total of over 41,000 tweets in this fashion. See Figure 7(A) for the diffusion network of this hashtag.

**@PeaceKaren\_25** This account did not disclose information about the identity of its owner, and generated a very large number of tweets (over 10,000 in four months). Almost all of these tweets supported several Republican candidates. Another account, @HopeMarie\_25, had a similar behavior to @PeaceKaren\_25 in retweeting the accounts of the same candidates and boosting the same Web sites. It did not produce any original tweets, and in addition it retweeted all of @PeaceKaren\_25's tweets, promoting that account. These accounts had also succeeded at creating a 'twitter bomb': for a time, Google searches for "gopleader" returned these tweets in the first page of results. A visualization of the interaction between these two accounts can be seen in Figure 7(B). Both accounts were suspended by Twitter by the time of this writing.

**gopleader.gov** This meme is the Web site of the Republican Leader John Boehner. It looks truthful because it is boosted by two suspicious accounts described above. The diffusion of this URL is shown in Figure 7(C).

### How Chris Coons budget works- uses tax \$ 2 attend dinners and fashion shows

This is one of a set of truthful memes smearing Chris Coons, the Democratic candidate for U.S. Senate from Delaware. Looking at the injection points of these memes, we uncovered a network of about ten bot accounts. They inject thousands of tweets with links to posts from the [freedomist.com](http://freedomist.com) Web site. To avoid detection by Twitter and increase visibility to different users, duplicate tweets are disguised by adding different hashtags and appending junk query parameters to the URLs. This works because many URL-shortening services ignore querystrings when processing redirect requests. To generate retweeting cascades, the bots also coordinate

nodes	Number of nodes
edges	Number of edges
mean_k	Mean degree
mean_s	Mean strength
mean_w	Mean edge weight in largest connected component
max_k(i, o)	Maximum (in,out)-degree
max_k(i, o)_user	User with max. (in,out)-degree
max_s(i, o)	Maximum (in,out)-strength
max_s(i, o)_user	User with max. (in,out)-strength
std_k(i, o)	Std. dev. of (in,out)-degree
std_s(i, o)	Std. dev. of (in,out)-strength
skew_k(i, o)	Skew of (in,out)-degree dist.
skew_s(i, o)	Skew of (in,out)-strength dist.
mean_cc	The mean size of connected components
max_cc	The size of the largest connected component
entry_nodes	Number of unique injections
num_truthy	Number of times 'truthy' button was clicked for the meme
sentiment scores	The six GPOMS sentiment dimensions

**Table 1: Features used in truthful classification.**

mentioning a few popular users. These targets get the appearance of receiving the same news from several different people, and are more likely to think it is true, and spread it to their followers. Most of the bot accounts in this network can be traced back to a single person who runs the [freedomist.com](http://freedomist.com) Web site. The diffusion network corresponding to this case is illustrated in Figure 7(D).

These are just a few instructive examples of characteristically truthful memes our system was able to identify. Two other networks of bots were shut down by Twitter after being detected by Truthy. In one case we observed the automated accounts using text segments drawn from newswire services to produce multiple legitimate-looking tweets in between the injection of URLs. These instances highlight several of the more general properties of truthful memes detected by our system.

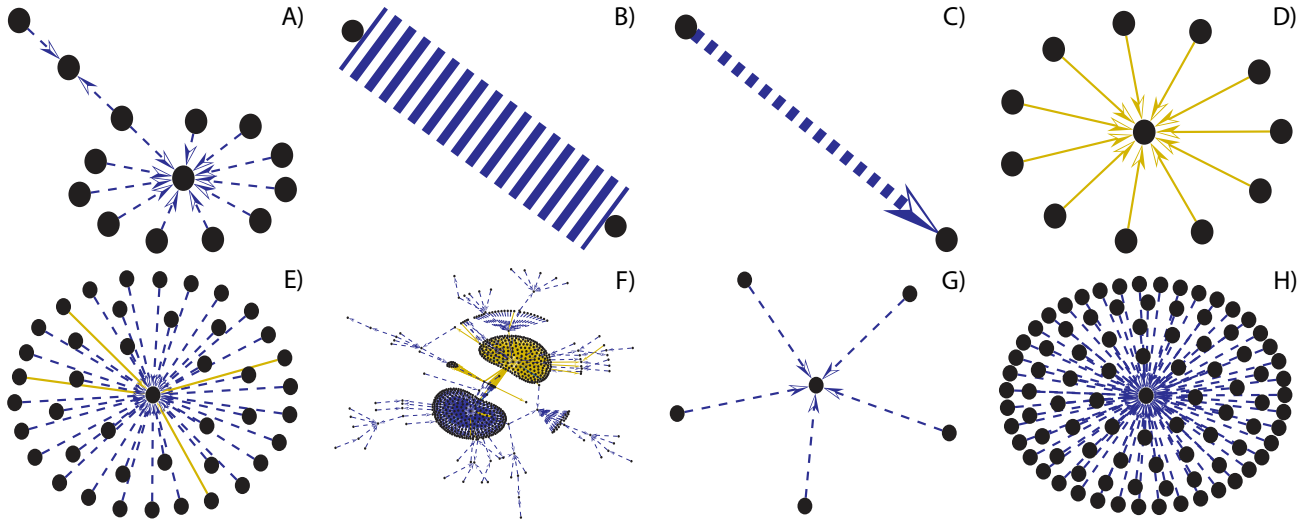
Figure 7 also shows the diffusion networks for four legitimate memes. One, #Truthy, was injected as an experiment by the NPR Science Friday radio program. Another, @senjohnmccain, displays two different communities in which the meme was propagated: one by retweets from @ladygaga in the context of discussion on the repeal of the "Don't ask, don't tell" policy on gays in the military, and the other by mentions of @senjohnmccain. A gallery with detailed explanations about various truthful and legitimate memes can be found on our Web site.<sup>2</sup>

## 6. TRUTHINESS CLASSIFICATION

As an application of the analyses performed by the Truthy system, we trained a binary classifier to automatically label legitimate and truthful memes.

We began by producing a hand-labeled corpus of training examples in three classes — 'truthy,' 'legitimate,' and 'remove.' We labeled these by presenting random memes to several human reviewers, and asking them to place each meme in one of the three categories. They were told to classify a meme as 'truthy' if a significant portion of the users involved in that meme appeared to be spreading it in misleading ways — e.g., if a number of the accounts tweeting about the meme appeared to be robots or sock puppets, the accounts appeared to follow only other propagators of the meme (clique behavior), or the users engaged in repeated reply/retweet

<sup>2</sup>[truthy.indiana.edu/gallery](http://truthy.indiana.edu/gallery)



**Figure 7: Diffusion networks of sample memes from our dataset.** Edges are represented using the same notation as in Figure 2. Four truthy memes are shown in the top row and four legitimate ones in the bottom row. (A) #ampat (B) @PeaceKaren\_25 (C) gopleader.gov (D) “How Chris Coons budget works- uses tax \$ 2 attend dinners and fashion shows” (E) #Truthy (F) @senjohnmccain (G) on.cnn.com/aVMu5y (H) “Obama said taxes have gone down during his administration. That’s ONE way to get rid of income tax — getting rid of income”

Classifier	Resampling?	Accuracy	AUC
AdaBoost	No	92.6%	0.91
AdaBoost	Yes	96.4%	0.99
SVM	No	88.3 %	0.77
SVM	Yes	95.6%	0.95

**Table 2: Performance of two classifiers with and without resampling training data to equalize class sizes. All results are averaged based on 10-fold cross-validation.**

	Positive	Negative	Positive	Negative
T	45 (12%)	294 (80%)	165 (45%)	188 (51%)
F	11 (3%)	16 (4%)	7 (2%)	6 (1%)
	No resampling		With resampling	

**Table 3: Confusion matrices for various classification strategies. Averages from 10-fold cross-validation are rounded to the nearest integers.**

exclusively with other users who had tweeted the meme. ‘Legitimate’ memes were described as memes representing normal, legitimate use of Twitter — several non-automated users conversing about a topic. The final category, ‘remove,’ was to be used for those memes that were in a foreign language, or otherwise did not seem to be related to American politics (#youth, for example). These memes were not used in the training or evaluation of classifiers.

After we had gathered several hundred annotations we observed an imbalance in our labeled data with less than 10% truthy. Rather than simply resampling, as is common practice in the case of class imbalance, we performed a second round of human annotations on previously-unlabeled memes predicted to be ‘truthy’ by the classifier trained in the previous round. In this way we were able to manually label a larger portion of truthy memes. Our final training dataset consisted of 366 training examples — 61 truthy memes and 305 legitimate ones. In those cases where multiple reviewers disagreed on the labeling of a meme, we determined the final label by a group discussion among all reviewers. The dataset is available online.<sup>3</sup>

We used the WEKA machine learning package [22] for classifier training, providing each classification strategy with 32 features about each meme, as shown in Table 1. We experimented with two classifiers: AdaBoost with DecisionStump, and SVM. As the number of instances of truthy memes was still less than instances of

legitimate ones, we also experimented with resampling the training data to balance the classes prior to classification. The performance of the classifiers is shown in Table 2, as evaluated by their accuracy and the area under their ROC curves. In all cases these preliminary results are quite encouraging, with accuracy around or above 90%. The best results are obtained by AdaBoost with resampling. Table 3 further shows the confusion matrices for AdaBoost. In this task, false negatives (truthy memes incorrectly classified as legitimate) are less desirable than false positives. In the worst case, the false negative rate is 5%. Table 4 shows the 10 most discriminative features, as determined by  $\chi^2$  analysis. Network features appear to be more discriminative than sentiment scores or the few user annotations that we collected.

## 7. DISCUSSION

We introduced a system for the real-time analysis of meme diffusion from microblog streams. The Klatsch framework will soon be released as open source. We described the Truthy system and Web site, which leverage this framework to track political memes in Twitter and help detect astroturfing campaigns in the context of U.S. political elections.

Our simple classification system yielded promising results in accurately detecting truthy memes based on features extracted from the topology of the diffusion networks. Using this system we have been able to identify a number of genuinely truthy memes. Though few of these exhibit the explosive growth characteristic of true vi-

<sup>3</sup>[cnets.indiana.edu/groups/nan/truthy](http://cnets.indiana.edu/groups/nan/truthy)



$\chi^2$	Rank	Feature
230 $\pm$ 4	1.0 $\pm$ 0.0	mean_w
204 $\pm$ 6	2.0 $\pm$ 0.0	mean_s
188 $\pm$ 4	4.3 $\pm$ 1.9	edges
185 $\pm$ 4	4.4 $\pm$ 1.1	skew_ko
183 $\pm$ 5	5.1 $\pm$ 1.3	std_si
184 $\pm$ 4	5.1 $\pm$ 0.9	skew_so
180 $\pm$ 4	6.7 $\pm$ 1.3	skew_si
177 $\pm$ 4	8.1 $\pm$ 1.0	max_cc
174 $\pm$ 4	9.6 $\pm$ 0.9	skew_ki
168 $\pm$ 5	11.5 $\pm$ 0.9	std_ko

**Table 4: Top 10 most discriminative features, according to a  $\chi^2$  analysis under 10-fold cross validation. Intervals represent the variation of the  $\chi^2$  or rank across the folds.**

ral memes, they are nonetheless clear examples of coordinated attempts to deceive Twitter users. Truthy memes are often spread initially by bots, causing them to exhibit pathological diffusion graphs relative to what is observed in the case of organic memes. These graphs can take many forms, including high numbers of unique injection points with few or no connected components, strong star-like topologies characterized by high average degree, and most tellingly large edge weights between dyads in graphs that exhibit either of the above properties.

The accuracy scores we observe in the classification task are quite high. We hypothesize that this performance is partially explained by the fact that a consistent proportion of the memes were failed attempts of starting a cascade. In these cases the networks reduced to isolated injection points or small components, resulting in trivial network features that allowed for easy classification.

Despite the fact that many of the memes discussed in this paper are characterized by small diffusion networks, it is important to note that this is the stage at which such attempts at deception must be identified. Once one of these attempts is successful at gaining the attention of the community, it will quickly become indistinguishable from an organic meme. Therefore, the early identification and termination of accounts associated with astroturf memes is critical.

In the future we intend to add more views to the website, including views on the users, such as the ages of the accounts, and tag clouds to interpret the sentiment analysis scores. We need to collect more labeled data about truthy memes in order to achieve more meaningful classification results, and will also explore the use of additional features in the classifiers, such as account ages for the most active users in a meme, and reputation features for users based on the memes to which they contribute. Another important area to address is that of sampling bias, since the properties of the sample made available in the Twitter gardenhose are currently unknown. To explore this, we intend to track injected memes of various sizes and with different topological properties of their diffusion graphs.

## Acknowledgments

We are grateful to Alessandro Vespignani, Ciro Catutto, Jose Ramasco, and Janette Lehmann for helpful discussions, Johan Bollen for his GPOMS code, Takis Metaxas and Eni Mustafaraj for inspiration and advice, and Ying Wang for Web design support. We thank the Gephi toolkit for aid in our visualizations and the many users who have provided feedback and annotations. We acknowledge support from NSF (grant No. IIS-0811994), Lilly Foundation (Data to Insight Center Research Grant), the Center for Complex Networks and Systems Research, and the School of Informatics and Computing at Indiana University, Bloomington.

## 8. REFERENCES

- [1] L. Adamic and N. Glance. The political blogosphere and the 2004 U.S. election: Divided they blog. In *LinkKDD '05: Proc. of the 3rd International Workshop on Link Discovery*, pages 36–43, 2005.
- [2] S. Aday, H. Farrel, M. Lynch, J. Sides, J. Kelly, and E. Zuckerman. Blogs and bullets: New media in contentious politics. Technical report, United States Institute of Peace, 2010.
- [3] Arbitron/Edison Internet and Multimedia Research Series. Twitter usage in America: 2010. Technical report, Edison Research, 2010.
- [4] S. Asur and B. A. Huberman. Predicting the future with social media. Technical Report arXiv:1003.5699, CoRR, 2010.
- [5] R. Axelrod. The dissemination of culture a model with local convergence and global polarization the dissemination of culture a model with local convergence and global polarization the dissemination of culture - a model with local convergence and global polarization. *J. Conflict Resolution*, 41:203, 1997.
- [6] A. Barrat, M. Barthelemy, and A. Vespignani. *Dynamical Processes on Complex Networks*. Cambridge University Press, 2008.
- [7] F. Benevenuto, G. Magno, T. Rodrigues, and V. Almeida. Detecting spammers on twitter. In *Proc. of the 7th Annual Collaboration, Electronic Messaging, Anti-Abuse and Spam Conf. (CEAS)*, 2010.
- [8] Y. Benkler. *The Wealth of Networks: How Social Production Transforms Markets and Freedom*. Yale University Press, 2006.
- [9] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. *J. Mach. Learn. Res.*, 3:993–1022, 2003.
- [10] J. Bollen, H. Mao, and A. Pepe. Determining the public mood state by analysis of microblogging posts. In *Proc. of the Alife XII Conf.* MIT Press, 2010.
- [11] J. Bollen, H. Mao, and X.-J. Zeng. Twitter mood predicts the stock market. Technical Report arXiv:1010.3003, CoRR, 2010.
- [12] D. Boyd, S. Golder, and G. Lotan. Tweet, tweet, retweet: Conversational aspects of retweeting on Twitter. In *43rd Hawaii International Conf. on System Sciences*, page 412, 2008.
- [13] C. Castellano, S. Fortunato, and V. Loreto. Statistical physics of social dynamics. *Rev. Mod. Phys.*, 81:591, 2009.
- [14] S. Deerwester and S. T. Dumais. Indexing by latent semantic analysis. *J. of the American Society for Information Science*, 1990.
- [15] N. A. Diakopoulos and D. A. Shamma. Characterizing debate performance via aggregated twitter sentiment. In *CHI '10: Proc. of the 28th International Conf. on Human Factors in Computing systems*, pages 1195–1198, New York, NY, USA, 2010. ACM.
- [16] P. Earle. Earthquake Twitter. *Nature Geoscience*, 3:221, 2010.
- [17] J. M. Epstein and R. L. Axtell. *Growing Artificial Societies: Social Science from the Bottom Up*. MIT Press, 1996.
- [18] W. Galuba, K. Aberer, D. Chakraborty, Z. Despotovic, and W. Kellerer. Outtweeting the Twitterers - Predicting Information Cascades in Microblogs. In *3rd Workshop on Online Social Networks (WOSN'10)*, 2010.

- [19] M. Gomez-Rodriguez, J. Leskovec, and A. Krause. Inferring networks of diffusion and influence. In *Proc. KDD*, pages 1019–1028, 2010.
- [20] B. Gonçalves, N. Perra, and A. Vespignani. The social brain online. In preparation, 2010.
- [21] C. Grier, K. Thomas, V. Paxson, and M. Zhang. @spam: the underground on 140 characters or less. In *Proc. of the 17th ACM Conf. on Computer and Communications Security (CCS)*, pages 27–37, 2010.
- [22] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten. The weka data mining software: An update. *SigKDD Explorations*, 11, 2009.
- [23] J. Heer and D. Boyd. Vizster: Visualizing online social networks. In *InfoVis 2005 IEEE Symposium on Information Visualization*, 2005.
- [24] C. Honeycutt and S. C. Herring. Beyond microblogging: Conversation and collaboration via Twitter. In *Proc. of the 42nd Hawaii International Conf. on System Sciences*, 2008.
- [25] J. Huang, K. M. Thornton, and E. N. Efthimiadis. Conversational tagging in Twitter. In *Proc. of the 21st ACM Conf. on Hypertext and Hypermedia*, 2010.
- [26] B. A. Huberman, D. M. Romero, and F. Wu. Social networks that matter: Twitter under the microscope. Technical Report arXiv:0812.1045, HP Labs, <http://arXiv.org/pdf/0812.1045> 2008.
- [27] T. N. Jagatic, N. A. Johnson, M. Jakobsson, and F. Menczer. Social phishing. *Communications of the ACM*, 50(10):94–100, 2007.
- [28] B. J. Jansen, M. Zhang, K. Sobel, and A. Chowdury. Twitter power: Tweets as electronic word of mouth. *J. of the American Society for Information Science*, 60:2169–2188, 2009.
- [29] A. Java, X. Song, T. Finin, and B. Tseng. Why we Twitter: understanding microblogging usage and communities. In *Proc. of the 9th WebKDD and 1st SNA-KDD 2007 workshop on Web mining and social network analysis*, 2007.
- [30] H. Kwak, C. Lee, H. Park, and S. Moon. What is Twitter, a social network or a news media? In *WWW '10: Proc. of the 19th international Conf. on World wide web*, pages 591–600, New York, NY, USA, 2010. ACM.
- [31] J. Leskovec, L. A. Adamic, and B. A. Huberman. Dynamics of viral marketing. In *ACM Transactions on the Web*, 2006.
- [32] J. Leskovec, L. Backstrom, and J. Kleinberg. Meme-tracking and the dynamics of the news cycle. In *KDD '09: Proc. of the 15th ACM SIGKDD international Conf. on Knowledge discovery and data mining*, page 497, 2009.
- [33] J. Leskovec and C. Faloutsos. Sampling from large graphs. In *KDD '06: Proc. of the 12th ACM SIGKDD international Conf. on Knowledge discovery and data mining*, pages 631–636, New York, NY, USA, 2006. ACM.
- [34] D. M. McNair, M. Lorr, and L. F. Droppelman. Profile of mood states. Technical report, Educational and Industrial Testing Service, 1971.
- [35] M. Mendoza, B. Poblete, and C. Castillo. Twitter under crisis: Can we trust what we RT? In *1st Workshop on Social Media Analytics (SOMA '10)*. ACM Press, July 2010.
- [36] E. Mustafaraj and P. Metaxas. From obscurity to prominence in minutes: Political speech and real-time search. In *WebSci10: Extending the Frontiers of Society On-Line*, page 317, 2010.
- [37] A. Pepe and J. Bollen. Between conjecture and memento: Shaping a collective emotional perception of the future. In *AAAI Spring Symposium on Emotion, Personality, and Social Behavior*, 2008.
- [38] A. Rapoport. Spread of information through a population with social structural bias: I. assumption of transitivity. *Bull. Math. Bio.*, 15:523, 1953.
- [39] S. Rasmussen and D. Schoen. *Mad as Hell: How the Tea Party Movement Is Fundamentally Remaking Our Two-Party System*. HarperCollins, 2010.
- [40] J. Ratkiewicz, F. Menczer, S. Fortunato, A. Flammini, and A. Vespignani. Traffic in Social Media II: Modeling Bursty Popularity. In *Proc. of the International Symposium on Social Intelligence and Networking (SIN-10)*. IEEE, 2010.
- [41] M. Raymond. How tweet it is!: Library acquires entire Twitter archive. <http://blogs.loc.gov/loc/2010/04/how-tweet-it-is-library-acquires-entire-twitter-archive/>.
- [42] D. M. Romero, W. Galuba, S. Asur, and B. A. Huberman. Influence and passivity in social media. Technical Report arXiv:1008.1253, CoRR, <http://arxiv.org/abs/1008.1253> 2010.
- [43] T. Sakaki, M. Okazaki, and Y. Matsuo. Earthquake shakes twitter users: real-time event detection by social sensors. In *WWW '10: Proc. of the 19th international Conf. on World wide web*, pages 851–860, New York, NY, USA, 2010. ACM.
- [44] J. Sankaranarayanan, H. Samet, B. E. Teitler, M. D. Lieberman, and J. Sperling. Twitterstand: news in tweets. In *GIS '09: Proc. of the 17th ACM SIGSPATIAL International Conf. on Advances in Geographic Information Systems*, pages 42–51, New York, NY, USA, 2009. ACM.
- [45] B. Suh, L. Hong, P. Pirolli, and E. H. Chi. Want to be retweeted? large scale analytics on factors impacting retweet in Twitter network. In *Proc. 2010 IEEE International Conf. on Social Computing*, 2010.
- [46] D. Tapscott. *Grown Up Digital: How the Net Generation is Changing Your World HC*. McGraw-Hill, 2008.
- [47] The Fox Nation. Stimulus \$ for coke monkeys. <http://politifi.com/news/Stimulus-for-Coke-Monkeys-267998.html>.
- [48] A. Tumasjan, T. O. Sprenger, P. G. Sandner, and I. M. Welpe. Predicting Elections with Twitter: What 140 Characters Reveal about Political Sentiment. In *Proc. of the Fourth International AAAI Conf. on Weblogs and Social Media (ICWSM)*, 2010.
- [49] A. H. Wang. Don't follow me: Twitter spam detection. In *Proc. 5th International Conf. on Security and Cryptography (SECRYPT)*, 2010.
- [50] H. Wang, W. Fan, P. S. Yu, and J. Han. Mining concept-drifting data streams using ensemble classifiers. In *KDD '03: Proc. of the ninth ACM SIGKDD international Conf. on Knowledge discovery and data mining*, pages 226–235, New York, NY, USA, 2003. ACM.
- [51] D. R. Wiese and B. E. Gronbeck. Campaign 2004: Developments in cyberpolitics. In R. E. Denton, editor, *The 2004 Presidential Campaign: A Communication Perspective*. Rowman & Littlefield, 2005.
- [52] S. Yardi, D. Romero, G. Schoenebeck, and danah boyd. Detecting spam in a Twitter network. *First Monday*, 15(1), 2009.