

OSoMe: The IUNI observatory on social media

Clayton A Davis^{1,2}, Giovanni Luca Ciampaglia^{Corresp., 1,3}, Luca Maria Aiello⁴, Keychul Chung², Michael D Conover⁵, Emilio Ferrara⁶, Alessandro Flammini^{1,2,3}, Geoffrey C Fox², Xiaoming Gao⁷, Bruno Gonçalves⁸, Przemyslaw A Grabowicz⁹, Kibeom Hong², Pik-Mai Hui^{1,2}, Scott McCaulay³, Karissa McKelvey¹⁰, Mark R Meiss¹¹, Snehal Patil¹², Chathuri Peli Kankanamalage³, Valentin Pentchev³, Judy Qiu², Jacob Ratkiewicz¹¹, Alex Rudnick¹¹, Benjamin Serrette³, Prashant Shiralkar^{1,2}, Onur Varol^{1,2}, Lilian Weng¹³, Tak-Lon Wu¹⁴, Andrew J Younge², Filippo Menczer^{1,2,3}

¹ Center for Complex Networks and Systems Research, Indiana University, Bloomington, United States

² School of Informatics and Computing, Indiana University, Bloomington, United States

³ Network Science Institute, Indiana University, Bloomington, United States

⁴ Bell Labs, London, United Kingdom

⁵ LinkedIn Inc., Mountain View, CA, United States

⁶ Information Sciences Institute, University of Southern California, Marina Del Rey, CA, United States

⁷ Facebook Inc., Boston, MA, United States

⁸ Center for Data Science, New York University, New York City, New York, United States

⁹ Max Planck Institute for Software Systems, Saarbrücken, Germany

¹⁰ U.S. Open Data, Oakland, CA, United States

¹¹ Google Inc., Mountain View, CA, United States

¹² Yahoo Inc., Sunnyvale, CA, United States

¹³ Affirm Inc., San Francisco, CA, United States

¹⁴ Amazon Inc., Seattle, WA, United States

Corresponding Author: Giovanni Luca Ciampaglia

Email address: gciampag@indiana.edu

The study of social phenomena is becoming increasingly reliant on big data from online social networks. Broad access to social media data, however, requires software development skills that not all researchers possess. Here we present the *IUNI Observatory on Social Media*, an open analytics platform designed to facilitate computational social science. The system leverages a historical, ongoing collection of over 70 billion public messages from Twitter. We illustrate a number of interactive open-source tools to retrieve, visualize, and analyze derived data from this collection. The Observatory, now available at osome.iuni.iu.edu, is the result of a large, six-year collaborative effort coordinated by the Indiana University Network Science Institute.

OSoMe: The IUNI Observatory on Social Media

Clayton A. Davis^{*1,2}, Giovanni Luca Ciampaglia^{1,3}, Luca Maria Aiello⁺⁴,
Keychul Chung², Michael Conover⁺⁵, Emilio Ferrara⁺⁶, Alessandro
Flammini^{1,2,3}, Geoffrey Fox², Xiaoming Gao⁺⁷, Bruno Gonçalves⁺⁸, Przemyslaw
Grabowicz⁺⁹, Alex Hong², Pik-Mai Hui^{1,2}, Scott McCaulay³, Karissa
McKelvey⁺¹⁰, Mark Meiss⁺¹¹, Snehal Patil⁺¹², Chathuri Peli Kankanamalage³,
Valentin Pentchev³, Judy Qiu², Jacob Ratkiewicz⁺¹¹, Alex Rudnick⁺¹¹, Benjamin
Serrette³, Prashant Shiralkar^{1,2}, Onur Varol^{1,2}, Lilian Weng⁺¹³, Tak-Lon Wu⁺¹⁴,
Andrew Younge², and Filippo Menczer^{1,2,3}

¹*Center for Complex Networks and Systems Research, Indiana University, USA*

²*School of Informatics and Computing, Indiana University, USA*

³*Network Science Institute, Indiana University, USA*

⁴*Bell Labs, UK*

⁵*LinkedIn Inc, USA*

⁶*Information Sciences Institute, University of Southern California, USA*

⁷*Facebook Inc, USA*

⁸*Center for Data Science, New York University, USA*

⁹*Max Planck Institute for Software Systems, Germany*

¹⁰*US Open Data, USA*

¹¹*Google Inc, USA*

¹²*Yahoo! Inc, USA*

¹³*Affirm Inc, USA*

¹⁴*Amazon Inc, USA*

*Corresponding author: claydavi@umail.iu.edu

[†]Work done at Indiana University.

Abstract

The study of social phenomena is becoming increasingly reliant on big data from on-line social networks. Broad access to social media data, however, requires software development skills that not all researchers possess. In this note we present the IUNI Observatory on Social Media, an open analytics platform designed to facilitate computational social science. The system leverages a historical corpus of over 70 billion public messages from Twitter for which data collection is ongoing. We illustrate a number of interactive publicly available tools to retrieve, visualize, and analyze derived data from this collection. Preliminary evaluation shows that system performance scales well with number and size of queries. The Observatory, now available at osome.iuni.iu.edu, is the result of a large, six-year collaborative effort coordinated by the Indiana University Network Science Institute.

13 Introduction

14 The collective processes of production, consumption, and diffusion of information on
 15 social media are starting to reveal a significant portion of human social life, yet scien-
 16 tists struggle to get access to data about it. Recent research has shown that social media
 17 can perform as ‘sensors’ for collective activity at multiple scales (Lazer et al., 2009). As
 18 a consequence, data extracted from social media platforms are increasingly used side-
 19 by-side with — and sometimes even replacing — traditional methods to investigate
 20 hard-pressing questions in the social, behavioral, and economic (SBE) sciences (King,
 21 2011; Moran et al., 2014; Einav and Levin, 2014). For example, interpersonal connections
 22 from Facebook have been used to replicate the famous experiment by Travers and Mil-
 23 gram (1969) on a global scale (Backstrom et al., 2012); the emotional content of social
 24 media streams has been used to estimate macroeconomic quantities in country-wide
 25 economies (Bollen et al., 2011; Choi and Varian, 2012; Antenucci et al., 2014); and im-
 26 agery from Instagram has been mined (De Choudhury et al., 2013; Andalibi et al., 2015)
 27 to understand the spread of depression among teenagers (Link et al., 1999).

28 A significant amount of work about information production, consumption, and dif-
 29 fusion has been thus aimed at modeling these processes and empirically discriminating
 30 among models of mechanisms driving the spread of memes on social media networks
 31 such as Twitter (Guille et al., 2013). A set of research questions relate to how social
 32 network structure, user interests, competition for finite attention, and other factors af-
 33 fect the manner in which information is disseminated and why some ideas cause viral
 34 explosions while others are quickly forgotten. Such questions have been addressed both
 35 in an empirical and in more theoretical terms.

36 Examples of empirical works concerned with these questions include geographic
 37 and temporal patterns in social movements (Conover et al., 2013b,a; Varol et al., 2014),
 38 the polarization of online political discourse (Conover et al., 2011b,a, 2012), the use of
 39 social media data to predict election outcomes (DiGrazia et al., 2013) and stock market
 40 movements (Bollen et al., 2011), the geographic diffusion of trending topics (Ferrara
 41 et al., 2013), and the lifecycle of information in the attention economy (Ciampaglia et al.,
 42 2015).

43 On the more theoretical side, agent-based models have been proposed to explain
 44 how limited individual attention affects what information we propagate (Weng et al.,
 45 2012), what social connections we make (Weng et al., 2013b), and how the structure
 46 of social and topical networks can help predict which memes are likely to become vi-
 47 ral (Weng et al., 2013a, 2014; Nematzadeh et al., 2014; Weng and Menczer, 2015).

48 Broad access by the research community to social media platforms is, however, lim-
 49 ited by a host of factors. One obvious limitation is due to the commercial nature of these
 50 services. On these platforms, data are collected as part of normal operations, but this is
 51 seldom done keeping in mind the needs of researchers. In some cases researchers have
 52 been allowed to harvest data through programmatic interfaces, or APIs. However, the
 53 information that a single researcher can gather through an API typically offers only a
 54 limited view of the phenomena under study; access to historical data is often restricted
 55 or unavailable (Zimmer, 2015). Moreover, these samples are often collected using ad-hoc
 56 procedures, and the statistical biases introduced by these practices are only starting to
 57 be understood (Morstatter et al., 2013; Ruths and Pfeffer, 2014; Hargittai, 2015).

58 A second limitation is related to the ease of use of APIs, which are usually meant
 59 for software developers, not researchers. While researchers in the SBE sciences are in-

creasingly acquiring software development skills (Terna, 1998; Raento et al., 2009; Healy and Moody, 2014), and intuitive user interfaces are becoming more ubiquitous, many tasks remain challenging enough to hinder research advances. This is especially true for those tasks related to the application of fast visualization techniques.

A third, important limitation is related to user privacy. Unfettered access to sensitive, private data about the choices, behaviors, and preferences of individuals is happening at an increasing rate (Tene and Polonetsky, 2012). Coupled with the possibility to manipulate the environment presented to users (Kramer et al., 2014), this has raised in more than one occasion deep ethical concerns in both the public and the scientific community (Kahn et al., 2014; Fiske and Hauser, 2014; Harriman and Patel, 2014; Vayena et al., 2015).

These limitations point to a critical need for opening social media platforms to researchers in ways that are both respectful of user privacy requirements and aware of the needs of SBE researchers. In the absence of such systems, SBE researchers will have to increasingly rely on closed or opaque data sources, making it more difficult to reproduce and replicate findings — a practice of increasing concern given recent findings about replicability in the SBE sciences (Open Science Collaboration, 2015).

Our long-term goal is to enable SBE researchers and the general public to openly access relevant social media data. As a concrete milestone of our project, here we present an *Observatory on Social Media* — an open infrastructure for sharing public data about information that is spread and collected through online social networks. Our initial focus has been on Twitter as a source of public microblogging posts. The infrastructure takes care of storing, indexing, and analyzing public collections and historical archives of big social data; it does so in an easy-to-use way, enabling broad access from scientists and other stakeholders, like journalists and the general public. We envision that data and analytics from social media will be integrated within a nation-wide network of social observatories. These data centers would allow access to a broad range of data about social, behavioral, and economic phenomena nationwide (King, 2011; Moran et al., 2014; Difranzo et al., 2014).

Our team has been working toward this vision since 2010, when we started collecting public tweets to visualize, analyze, and model meme diffusion networks. The IUNI Observatory on Social Media (OSoMe) presented here was launched in early May 2016. It was developed through a collaboration between the Indiana University Network Science Institute (IUNI, iuni.iu.edu), the IU School of Informatics and Computing (SoIC, soic.indiana.edu), and the Center for Complex Networks and Systems Research (CNetS, cnets.indiana.edu). It is available at osome.iuni.iu.edu.¹

Data Source

Social media data possess unique characteristics. Besides rich textual content, explicit information about the originating social context is generally available. Information often includes timestamps, geolocations, and interpersonal ties. The Twitter dataset is

¹The OSoMe website is also available at truthy.indiana.edu. The original website was created to host our first demo, motivated by the application of social media analytics to the study of “astroturf,” or artificial grassroots social media campaigns orchestrated through fake accounts and social bots (Ratkiewicz et al., 2011b). The *Truthy* nickname was later adopted in the media to refer to the entire project. The current website includes information about other research projects on information diffusion and social bots from our lab.




Figure 1: Number of monthly messages collected and indexed by OSoMe. System failures have caused occasional interruptions of the collection system.

a prototypical example (McKelvey and Menczer, 2013b,a). The Observatory on Social Media is built around a Terabyte-scale historical (and ongoing) collection of tweets. The data source is a large sample of public posts made available by Twitter through elevated access to their streaming API, granted to a number of academic labs. All of the tweets from this sample are stored, resulting in a corpus of approximately **70 billion public tweets** dating back to mid-2010.²

An important caveat about the use of these data for research is that possible sampling biases are unknown. When Twitter first made this stream available to the research community, it indicated that the stream contains a random 10% sample of all public tweets. However, no further information about the sampling method was disclosed. Other streaming APIs have been shown to provide a non-uniform sample (Morstatter et al., 2013). Even assuming that tweets are randomly sampled, it should be noted that the collection does not automatically translate into a representative sample of the underlying population of Twitter user, or of the topics discussed. This is because the distribution of activity is highly skewed across users and topics (Weng et al., 2012) and, as a result, active users and popular topics are better represented in the sample. Sampling biases may also have evolved over time. For example, the fraction of public tweets with exact geolocation coordinates has decreased from approximately 3% in the past to approximately 0.3% due to the recent change of location privacy settings in Twitter mobile clients from “opt-out” to “opt-in.” This change was motivated by public privacy concerns about location tracking. This and other platform changes may significantly affect the composition of our sample in ways that we are unable to assess.

The high-speed stream from which the data originates has a rate that ranges in the order of $10^6 - 10^8$ tweets/day. Figure 1 illustrates the growth of the Twitter collection over time.

²Research based on this data was deemed exempt from review by the Indiana University IRB under Protocol #1102004860.

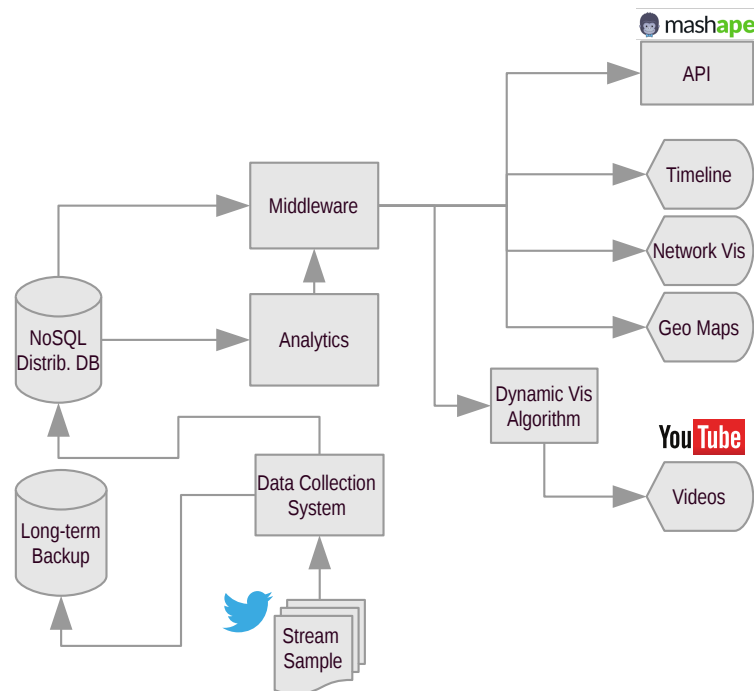


Figure 2: Flowchart diagram of the OSoMe architecture. Arrows indicate flow of data.

System Architecture

Performing analytics at this scale presents specific challenges. The most obvious has to do with the design of a suitable architecture for processing such a large volume of data. This requires a scalable storage substrate and efficient query mechanisms.

The core of our system is based on a distributed storage cluster composed of 10 compute nodes, each with $12 \times 3\text{TB}$ disk drives, $2 \times 146\text{GB}$ RAID-1 drives for the operative system, 64GB RAM, and $2 \times \text{Xeon 2650}$ CPUs with 8 cores each (16 total per node). Access to the nodes is provided via two head nodes, each equipped with 64GB RAM, and $4 \times \text{Xeon 2650}$ CPUs with four cores each (24 total per node), using 1GB ethernet infiniband.

The software architecture the Observatory builds upon the Apache Big Data Stack (ABDS) framework (Jha et al., 2014; Qiu et al., 2014; Fox et al., 2014). Development has been driven over the years by the need for increasingly demanding social media analytics applications (Gao et al., 2011; Gao and Qiu, 2013, 2014; Gao et al., 2014, 2015; Wu et al., 2016). A key idea behind our enhancement of the ABDS architecture is the shift from standalone systems to modules; multiple modules can be used within existing software ecosystems. In particular, we have focused our efforts on enhancing two well-known Apache modules, Hadoop (The Apache Software Foundation, 2016b) and HBase (The Apache Software Foundation, 2016a).

The architecture is illustrated in Figure 2. The *data collection system* receives data from the Twitter Streaming API. Data are first stored on a temporary location and then loaded into the distributed storage layer on a daily basis. At the same time, *long-term backups* are stored on tape to allow recovery in case of data loss or catastrophic events.

The design of the *NoSQL distributed DB* module was guided by the observation that

queries of social media data often involve unique constraints on the textual and social context such as temporal or network information. To address this issue, we leveraged the HBase system as the storage substrate and extended it with a flexible indexing framework. The resulting *IndexedHBase* module allows one to define fully customizable text index structures that are not supported by current state-of-the-art text indexing systems, such as Solr (The Apache Software Foundation, 2016c). The custom index structures can embed contextual information necessary for efficient query evaluation. The *IndexedHBase* software is publicly available (Wiggins et al., 2016).

The pipelines commonly used for social media data analysis consist of multiple algorithms with varying computation and communication patterns. For example, building the network of retweets of a given hashtag will take more time and computational resources than just counting the number of posts containing the hashtag. Moreover, the temporal resolution and aggregation windows of the data could vary dramatically, from seconds to years. A number of different processing frameworks could be needed to perform such a wide range of tasks. To design the *analytics* module of the Observatory we choose Hadoop, a standard framework for Big Data analytics. We use YARN (The Apache Software Foundation, 2016d) to achieve efficient execution of the whole pipeline, and integrate it with *IndexedHBase*. An advantage deriving from this choice is that the overall software stack can dynamically adopt different processing frameworks to complete heterogeneous tasks of variable size.

A distributed task queue, and an in-memory key/value store implement the *middleware* layer needed to submit queries to the backend of the Observatory. We use Celery (Solem and Contributors, 2016) and Redis (Sanfilippo, 2016) to implement such layer. The task queue limits the number of concurrent jobs processed by the system according to the task type (index-only vs map/reduce) to prevent extreme degradation of performance due to very high load.

The Observatory user interface follows a modular architecture too, and is based on a number of apps, which we describe in greater detail in the following section. Three of the apps (*Timeline*, *Network visualization*, and *Geographic maps*) are directly accessible within OSoMe through Web interfaces. We rely on the popular video-sharing service YouTube for the fourth app, which generates meme diffusion movies (*Videos*). The movies are rendered using a fast *dynamic visualization algorithm* that we specifically designed for temporal networks. The algorithm captures only the most persistent trends in the temporal evolution, at the expense of high-frequency churn (Grabowicz et al., 2014). The software is freely available (Grabowicz and Aiello, 2013). Finally, the Observatory provides access to raw data via a programmatic interface (*API*).

Applications

Storing and indexing tens of billions of tweets is of course pointless without a way to make sense of such a huge trove of information. The Observatory lowers the barrier of entry to social media analysis by providing users with several ready-to-use, Web-based data visualization tools. Visualization techniques allow users to make sense of complex data and patterns (Card, 2009), and let them explore the data and try different visualization parameters (Rafaeli, 1988). In the following, we give a brief overview of the available tools.

It is important to note that, in compliance with the Twitter terms of service (Twitter,




Figure 3: Number of tweets per day about the Super Bowl (in blue) and the World Series (in orange), from September 2015 through February 2016. The Y-axis is in logarithmic scale, shifted by one to account for null counts. The plot shows two outages in the data collection that occurred around mid-November 2015 and mid-January 2016.

194 Inc., 2016), OSoMe does not provide access to the content of tweets, nor of Twitter user
 195 objects. However, researchers can obtain numeric object identifiers of tweets in response
 196 to their queries. This information can then be used to retrieve tweet content via the
 197 official Twitter API. (There is one exception, described below.) Another necessary step
 198 to comply with the Twitter terms is to remove deleted tweets from the database. Using
 199 a Redis queue, we collect deletion notices from the public Twitter stream, and then feed
 200 them to a backend task for deletion.

201 Temporal Trends

202 The *Trends* tool produces time series plots of the number of tweets including one or
 203 more given hashtags; it can be compared to the service provided by Google Trends,
 204 which allows users to examine the interest toward a topic reflected by the volume of
 205 search queries submitted to Google over time.

206 Users may specify multiple terms in one query, in which case all tweets containing
 207 any of the terms will be computed; and they can perform multiple queries, to allow
 208 comparisons between different topics. For example, let us compare the relative tweet
 209 volumes about the World Series and the Superbowl. We want our Super Bowl timeline
 210 to count tweets containing any of #SuperBowl, #SuperBowl50, or #SB50. Since hashtags
 211 are case-insensitive and we allow trailing wildcards, this query would be “#superbowl*,
 212 #sb50.” Adding a timeline for the “#worldseries” query results in the plot seen in Fig-
 213 ure 3. Each query on the Trends tool takes 5–10 seconds; this makes the tool especially
 214 suitable for interactive exploration of Twitter conversation topics.

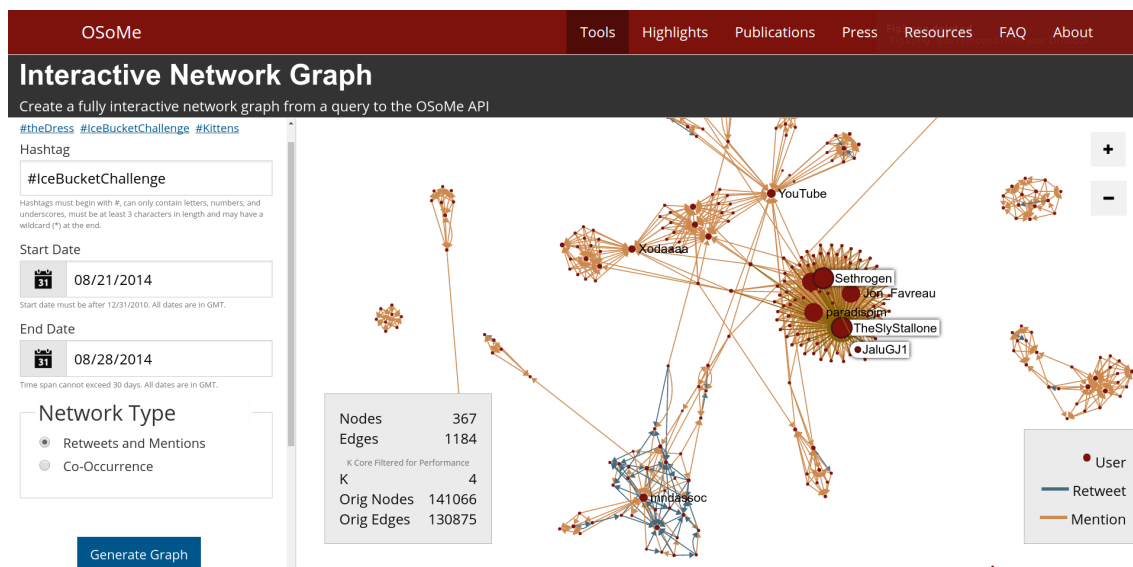


Figure 4: Interactive Network Visualization Tool. A detail of the network of retweets and mention for a hashtag commonly linked to “Ice Bucket Challenge,” a popular Internet phenomenon from 2014. The size of a node is proportional to its strength (weighted degree). The detail shows the patterns of mention and information broadcasting occurring between celebrities, as the viral challenge was taking off.

Diffusion and Co-occurrence Networks

In a diffusion network, nodes represent users and an edge drawn between any two nodes indicates an exchange of information between those two users. For example, a user could rebroadcast (*retweet*) the status of another user to her followers, or she could address another user in one of her statuses by including a mention to their user handle (*mention*). Edges have a weight to represent the number of messages connecting two nodes. They may also have an intrinsic direction to represent the flow of information. For example, in the retweet network for the hashtag #IceBucketChallenge, an edge from user i to user j indicates that j retweeted tweets by i containing the hashtag #IceBucketChallenge. Similarly, in a mention network, an edge from i to j indicates that i mentioned j in tweets containing the hashtag. Information diffusion network, sometimes also called information cascades, have been the subject of intense study in recent years (Gruhl et al., 2004; Weng et al., 2012; Bakshy et al., 2012; Weng et al., 2013b,a; Romero et al., 2011).

Another type of network visualizes how hashtags co-occur with each other. Co-occurrence networks are also weighted, but undirected: nodes represent hashtags, and the weight of an edge between two nodes is the number of tweets containing both of those hashtags.

OSoMe provides two tools that allow users to explore diffusion and and co-occurrence networks.

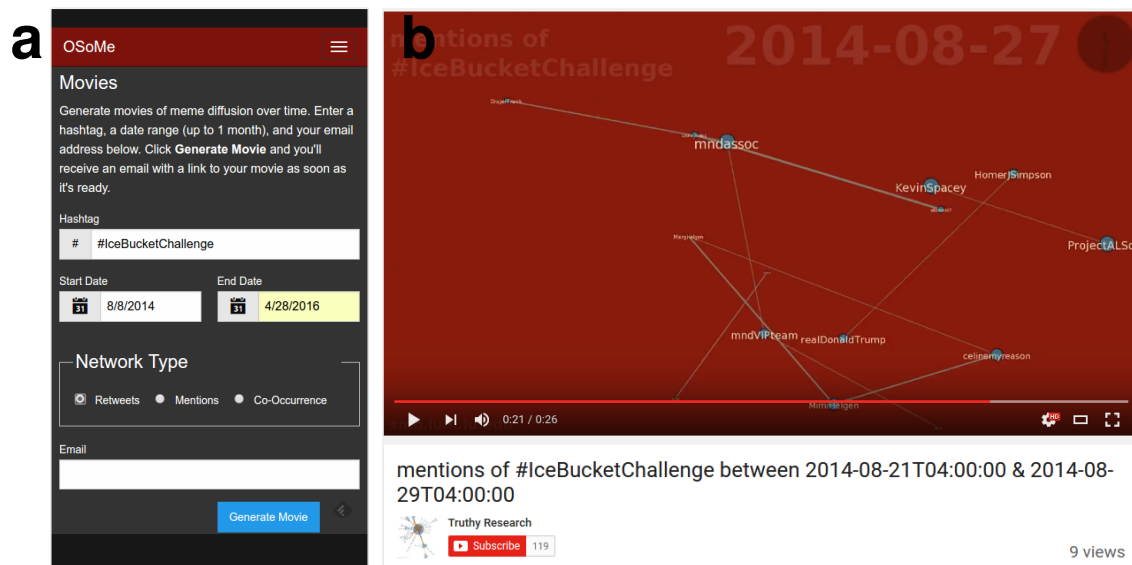


Figure 5: Temporal information diffusion movies. (a) The interface of the *Movies* tool let users specify a hashtag, a temporal interval, and the type of diffusion ties to visualize (retweets, mentions, or hashtag co-occurrence). (b) Example of a generated movie frame, showing a retweet network for the #IceBucketChallenge hashtag.

Interactive Network Visualization

The *Networks* tool enables the visualization of how a given hashtag spreads through the social network via retweets and mentions (Figure 4) or what hashtags co-occur with a given hashtag. The resulting network diagrams, created using a force-directed layout (Kamada and Kawai, 1989), can reveal topological patterns such as influential or highly-connected users and tightly-knit communities. Users can click on the nodes and edges to find out more information about the entities displayed — users, tweets, retweets, and mentions — directly from Twitter. Network are cached to enable fast access to previously-created visualizations.

For visualization purposes, the size of large networks is reduced by extracting their k -core (Alvarez-Hamelin et al., 2005) with k sufficiently large to display 1,000 nodes or less ($k = 5$ in the example of Figure 4). The use of this type of filter implies a bias toward densely connected portions of diffusion networks, where most of the activity occurs, and toward the most influential participants.

The tool allows access to Twitter content, such as hashtags and user screen names. This content is available both through the interactive interface itself, and as a downloadable JSON file. To comply with the Twitter terms, the k -core filter also limits the number of edges (tweets).

Animations

Because tweet data are time resolved, the evolution of a diffusion or co-occurrence network can be also visualized over time. Currently the *Networks* tool visualizes only static networks aggregated over the entire search period specified by the user; we aim to add the ability to observe the network evolution over time, but in the meantime we also pro-

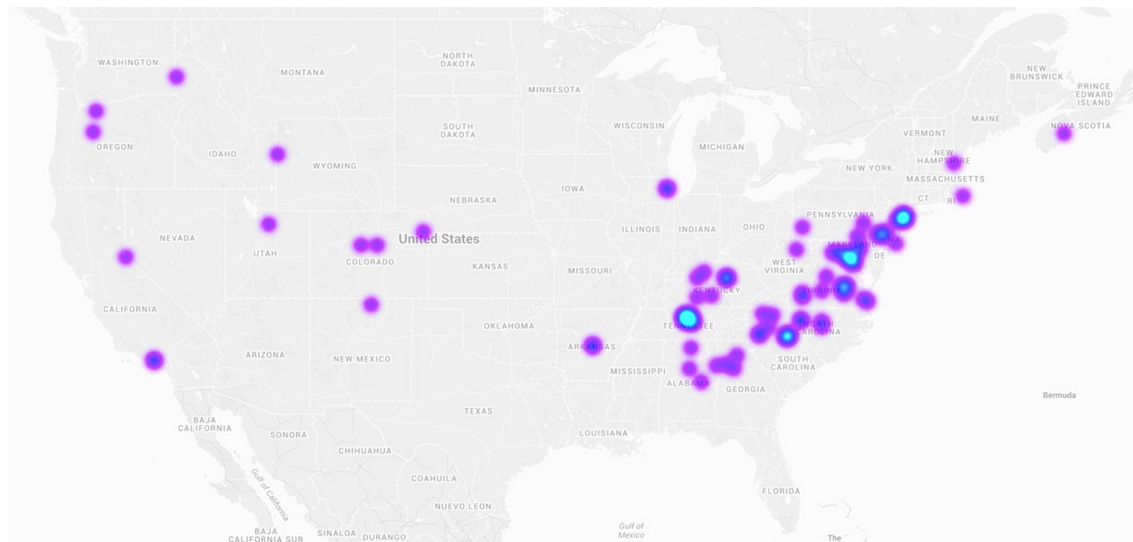


Figure 6: Heatmap of tweets containing the hashtag #snow on January 22, 2016, the day of a large snowstorm over the Eastern United States.

vide the *Movies* tools, an alternative service that lets users generate animations of such processes (Figure 5). We have successfully experimented with fast visualization techniques in the past, and have found that edge filtering is the best approach for efficiently visualizing networks that undergo a rapid churn of both edges and nodes. We have therefore deployed a fast filtering algorithm developed by our team (Grabowicz et al., 2014). The user-generated videos are uploaded to YouTube, and we cache the videos in case multiple users try to visualize the same network.

Geographic maps

Online social networks are implicitly embedded in space, and the spatial patterns of information spread have started to be investigated in recent years (Ferrara et al., 2013; Conover et al., 2013a). The *Maps* tool enables the exploration of information diffusion through geographic space and time. A subset of tweets contain exact latitude/longitude coordinates in their metadata. By aggregating these coordinates into a heatmap layer superimposed on a world map, one can observe the geographic signature of the attention being paid to a given meme. Figure 6 shows an example. Our online tool goes one step further, allowing the user to explore how this geographic signature evolves over a specified time period, via a slider widget.

It takes at least 30 seconds to prepare one of these visualizations *ex novo*. We hope to reduce this lead time with some backend indexing improvements. To enable exploration, we cache all created heatmaps for a period of one week. While cached, the heatmaps can be retrieved instantly, enabling other users to browse and interact with these previously-created visualizations. In the future we hope to experiment with overlaying diffusion networks on top of geographical maps, for example using multi-scale backbone extraction (Serrano et al., 2009) and edge bundling techniques (Selassie et al., 2011).

An important caveat for the use of the maps tool is that it is based on the very small

percentage of tweets that contain exact geolocation coordinates. Furthermore, as already discussed, this percentage has changed over time.

API

We expect that the majority of users of the Observatory will interact with its data primarily through the tools described above. However, since more advanced data needs are to be expected, we also provide a way to export the data for those who wish to create their own visualizations and develop custom analyses. This is possible either within the tools, via export buttons, and through a read-only HTTP API.

The OSoMe API is deployed via the Mashape management service. Four public methods are currently available. Each takes as input a time interval and a list of tokens (hashtags and/or usernames):

- `tweet-id`: returns a list of tweet IDs mentioning at least one of the inputs in the given interval;
- `counts`: returns a count of the number of tweets mentioning each input token in the given interval;
- `time-series`: for each day in the given time interval, returns a count of tweets matching any of the input tokens;
- `user-post-count`: returns a list of user IDs mentioning any of the tokens in the given time frame, along with a count of matching tweets produced by each user.

Evaluation

In the first several weeks since launch, the OSoMe infrastructure has served a large number of requests, as shown in Figure 7. The spike corresponds to May 6, 2016, the date of a press release about the launch. Most of these requests complete successfully, with no particular deterioration for increasing loads (Figure 8).

To evaluate the scalability of the Hadoop-based analytics tools with increasing data size, we plot in Figure 9 the run time of queries submitted by users through OSoMe interactive tools, as a function of the number of tweets matching the query parameters. We observe a sublinear growth, suggesting that the system scales well with job size. A job may take from approximately 30 seconds to several hours depending on many factors such as system load and number of tweets processed. However, even different queries that process the same number of tweets may perform differently, depending on the width of the query time window. This is partly due to “hotspotting”: the temporal locality of our data layout across the nodes of the storage cluster causes decreases in performance when different Hadoop mappers access the same disks. A query spanning a short period of time runs slower than one matching the same number of tweets over a longer period. These results suggest that our data layout design may need to be reconsidered in future development.

Finally, we tested the scalability of queries using the HBase index with the load. Figure 10 shows that the total run time is not strongly affected by the number of concurrent jobs, up to the size of the task queue (32). For larger loads, run time scales linearly as expected.




Figure 7: Number of daily requests to the Observatory, including both API calls and interactive queries. The inset shows the same data, on a logarithmic scale.




Figure 8: Fraction of successful requests as a function of daily request load. Error bars are standard errors within logarithmic bins.




Figure 9: Run time versus number of tweets processed by Hadoop-based interactive tools. The line is a guide for the eye corresponding to linear growth.




Figure 10: Run time versus number of concurrent jobs that use the HBase index. The line is a guide for the eye representing linear growth.

Conclusion

The IUNI Observatory on Social Media is the culmination of a large collaborative effort at Indiana University that took place over the course of six years. We hope that it will facilitate computational social science and make big social data easier to analyze by a broad community of researchers, reporters, and the general public. The lessons learned during the development of the infrastructure may be helpful for future endeavors to foster data-intensive research in the social, behavioral, and economic sciences.

We welcome feedback from researchers and other end users about usability and usefulness of the tools presented here. In the future, we plan to carry out user studies and tutorial workshops to gain feedback on effectiveness of the user interfaces, efficiency of the tools, and desirable extensions.

We encourage the research community to create new social media analytic tools by building upon our system. As an illustration, we created a mashup of the OSoMe API with the BotOrNot API (Davis et al., 2016), also developed by our team, to evaluate the extent to which Twitter campaigns are sustained by social bots. The software is freely available online (Davis, 2016).

The opportunities that arise from the Observatory, and from computational social science in general, could have broad societal impact. Systematic attempts to mislead the public on a large scale through “astroturf” campaigns and social bots have been uncovered using big social data analytics, inspiring the development of machine learning methods to detect these abuses (Ratkiewicz et al., 2011a; Ferrara et al., 2016; Subrahmanian et al., 2016). Allowing citizens to observe how memes spread online may help raise public awareness of the potential dangers of social media manipulation.

Acknowledgements

The authors would like to acknowledge Alessandro Vespignani and Johan Bollen for discussions leading to the early vision of an Observatory on Social Media; and Rob Henderson, Shing-Shong (Bruce) Shei, Gary Miksik, Allan Streib, and Koji Tanaka for their kind assistance with system administration. We are deeply grateful to Twitter for supporting computational social science research, including the efforts described in this paper, by granting our lab elevated access to the public stream of tweets.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the funding agencies.

References

- J. I. Alvarez-Hamelin, L. Dall’Asta, A. Barrat, and A. Vespignani. Large scale networks fingerprinting and visualization using the k-core decomposition. In *Advances in Neural Information Processing Systems 18 (NIPS)*, pages 41–50, 2005.
- N. Andalibi, P. Ozturk, and A. Forte. Depression-related imagery on instagram. In *Proc. 18th ACM Conf. Companion on Computer Supported Cooperative Work & Social Computing (CSCW)*, pages 231–234, 2015. doi: 10.1145/2685553.2699014.

- 365 D. Antenucci, M. Cafarella, M. Levenstein, C. Ré, and M. D. Shapiro. Using social media
366 to measure labor market flows. Working Paper 20010, National Bureau of Economic
367 Research, March 2014.
- 368 L. Backstrom, P. Boldi, M. Rosa, J. Ugander, and S. Vigna. Four degrees of separation.
369 In *Proceedings of the 4th Annual ACM Web Science Conference, WebSci '12*, pages 33–42,
370 2012. doi: 10.1145/2380718.2380723.
- 371 E. Bakshy, I. Rosenn, C. Marlow, and L. Adamic. The role of social networks in infor-
372 mation diffusion. In *Proceedings of the 21st ACM International Conference on World Wide*
373 *Web*, pages 519–528, 2012.
- 374 J. Bollen, H. Mao, and X. Zeng. Twitter mood predicts the stock market. *Journal of*
375 *Computational Science*, 2(1):1–8, 2011.
- 376 S. Card. Information visualization. In *Human-computer interaction: design issues, solutions,*
377 *and applications*, pages 181–216. CRC Press, 2009.
- 378 H. Choi and H. Varian. Predicting the present with google trends. *Economic Record*, 88
379 (s1):2–9, 2012.
- 380 G. L. Ciampaglia, A. Flammini, and F. Menczer. The production of information in the
381 attention economy. *Scientific Reports*, 5:9452, 2015. doi: 10.1038/srep09452.
- 382 M. Conover, B. Gonçalves, J. Ratkiewicz, A. Flammini, and F. Menczer. Predicting the
383 political alignment of Twitter users. In *Proc. 3rd IEEE Conference on Social Computing*
384 *(SocialCom)*, 2011a.
- 385 M. Conover, J. Ratkiewicz, M. Francisco, B. Gonçalves, A. Flammini, and F. Menczer.
386 Political polarization on Twitter. In *Proc. 5th International AAAI Conference on Weblogs*
387 *and Social Media (ICWSM)*, 2011b.
- 388 M. Conover, C. Davis, E. Ferrara, K. McKelvey, F. Menczer, and A. Flammini. The
389 geospatial characteristics of social movement communication networks. *PLoS ONE*, 8
390 (3):e55957, 2013a.
- 391 M. Conover, E. Ferrara, F. Menczer, and A. Flammini. The digital evolution of occupy
392 wall street. *PLoS ONE*, 8(3):e64679, 2013b.
- 393 M. D. Conover, B. Gonçalves, A. Flammini, and F. Menczer. Partisan asymmetries in
394 online political activity. *EPJ Data Science*, 1:6, 2012.
- 395 C. A. Davis. OSoMe Mashups, 2016. URL [https://github.com/IUNetSci/](https://github.com/IUNetSci/osome-mashups/)
396 [osome-mashups/](https://github.com/IUNetSci/osome-mashups/). Last accessed July 29, 2016.
- 397 C. A. Davis, O. Varol, E. Ferrara, A. Flammini, and F. Menczer. Botornot: A system
398 to evaluate social bots. In *Proc. WWW Developers Day Workshop*, 2016. doi: 10.1145/
399 2872518.2889302. URL <http://arxiv.org/abs/1602.00975>.
- 400 M. De Choudhury, M. Gamon, S. Counts, and E. Horvitz. Predicting depression via
401 social media. In *Proc. 7th Intl. AAAI Conf. on Weblogs and Social Media (ICWSM)*, 2013.

- 402 D. Difranzo, J. S. Erickson, M. J. K. T. Gloria, J. S. Luciano, D. L. McGuinness, and
403 J. Hendler. The web observatory extension: Facilitating web science collaboration
404 through semantic markup. In *Proc. 23rd Intl. Conf. on World Wide Web Companion*,
405 pages 475–480, 2014. doi: 10.1145/2567948.2576936.
- 406 J. DiGrazia, K. McKelvey, J. Bollen, and F. Rojas. More tweets, more votes: Social media
407 as a quantitative indicator of political behavior. *PLoS ONE*, 8(11), 2013.
- 408 L. Einav and J. Levin. Economics in the age of big data. *Science*, 346(6210):
409 1243089–1243089, Nov 2014. ISSN 1095-9203. doi: 10.1126/science.1243089.
- 410 E. Ferrara, O. Varol, F. Menczer, and A. Flammini. Traveling Trends: Social Butterflies
411 or Frequent Fliers? In *Proc. 1st ACM Conf. on Online Social Networks (COSN)*, pages
412 213–222, 2013.
- 413 E. Ferrara, O. Varol, C. Davis, F. Menczer, and A. Flammini. The rise of social bots.
414 *Comm. ACM*, 57(7):96–104, 2016. doi: 10.1145/2818717.
- 415 S. T. Fiske and R. M. Hauser. Protecting human research participants in the age of big
416 data. *Proceedings of the National Academy of Sciences*, 111(38):13675–13676, 2014. doi:
417 10.1073/pnas.1414626111.
- 418 G. C. Fox, S. Jha, J. Qiu, and A. Luckow. Towards an understanding of facets and
419 exemplars of big data applications. In *Proceedings of 20 Years of Beowulf: Workshop to*
420 *Honor Thomas Sterling’s 65th Birthday*, pages 7–16, 2014.
- 421 X. Gao and J. Qiu. Supporting end-to-end social media data analysis with the Indexed-
422 HBase platform. In *Proceedings of the 6th Workshop on Many-Task Computing on Clouds,*
423 *Grids, and Supercomputers (MTAGS) at SC13*, 2013.
- 424 X. Gao and J. Qiu. Supporting queries and analyses of large-scale social media data with
425 customizable and scalable indexing techniques over nosql databases. In *Proceedings*
426 *of the 14th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing*
427 *(CCGrid 2014)*, pages 587–590, 2014.
- 428 X. Gao, V. Nachankar, and J. Qiu. Experimenting Lucene Index on HBase in an HPC En-
429 vironment. In *Proceedings of ACM High Performance Computing meets Databases workshop*
430 *(HPCDB’11) at SuperComputing 11*, pages 25–28, 2011.
- 431 X. Gao, E. Roth, K. McKelvey, C. Davis, A. Younge, E. Ferrara, F. Menczer, and J. Qiu.
432 Supporting a social media observatory with customizable index structures: Archi-
433 tecture and performance. In *Cloud Computing for Data Intensive Applications*, pages
434 401–427. Springer, 2014.
- 435 X. Gao, E. Ferrara, and J. Qiu. Parallel clustering of high-dimensional social media data
436 streams. In *Proc. 15th IEEE/ACM International Symposium on Cluster, Cloud and Grid*
437 *Computing (CCGrid)*, pages 323–332, 2015.
- 438 P. A. Grabowicz and L. M. Aiello. Fastviz, 2013. URL [https://github.com/WICI/](https://github.com/WICI/fastviz)
439 [fastviz](https://github.com/WICI/fastviz). Last accessed July 21, 2016.

- 440 P. A. Grabowicz, L. M. Aiello, and F. Menczer. Fast filtering and animation of
441 large dynamic networks. *EPJ Data Science*, 3(1):27, 2014. doi: 10.1140/epjds/
442 s13688-014-0027-8.
- 443 D. Gruhl, R. Guha, D. Liben-Nowell, and A. Tomkins. Information diffusion through
444 blogspace. In *Proceedings of the 13th International ACM Conference on World Wide Web*,
445 WWW '04, pages 491–501, 2004. doi: 10.1145/988672.988739.
- 446 A. Guille, H. Hacid, C. Favre, and D. A. Zighed. Information diffusion in online social
447 networks. *SIGMOD Rec.*, 42(1):17, 2013. doi: 10.1145/2503792.2503797.
- 448 E. Hargittai. Is Bigger Always Better? Potential Biases of Big Data Derived from Social
449 Network Sites. *The Annals of the American Academy of Political and Social Science*, 659(1):
450 63—76, 2015. doi: 10.1177/0002716215570866.
- 451 S. Harriman and J. Patel. The ethics and editorial challenges of internet-based research.
452 *BMC Med*, 12(1), 2014. doi: 10.1186/s12916-014-0124-3.
- 453 K. Healy and J. Moody. Data visualization in sociology. *Annual review of sociology*, 40:
454 105—128, 2014. doi: 10.1146/annurev-soc-071312-145551.
- 455 S. Jha, J. Qiu, A. Luckow, P. Mantha, and G. C. Fox. A tale of two data-intensive
456 paradigms: Applications, abstractions, and architectures. In *Proceedings of the 3rd*
457 *International Congress on Big Data Conference (IEEE BigData)*, 2014.
- 458 J. P. Kahn, E. Vayena, and A. C. Mastroianni. Opinion: Learning as we go: Lessons from
459 the publication of facebook’s social-computing research. *Proceedings of the National*
460 *Academy of Sciences*, 111(38):13677–13679, 2014. doi: 10.1073/pnas.1416405111.
- 461 T. Kamada and S. Kawai. An algorithm for drawing general undirected graphs. *Infor-*
462 *mation Processing Letters*, 31(1):7 – 15, 1989.
- 463 G. King. Ensuring the data-rich future of the social sciences. *Science*, 331(6018):719–721,
464 2011. doi: 10.1126/science.1197872.
- 465 A. D. Kramer, J. E. Guillory, and J. T. Hancock. Experimental evidence of massive-scale
466 emotional contagion through social networks. *Proceedings of the National Academy of*
467 *Sciences*, 111(24):8788–8790, 2014.
- 468 D. Lazer, A. Pentland, L. Adamic, S. Aral, A.-L. Barabási, D. Brewer, N. Christakis,
469 N. Contractor, J. Fowler, M. Gutmann, T. Jebara, G. King, M. Macy, D. Roy, and
470 M. Van Alstyne. Computational social science. *Science*, 323(5915):721–723, 2009. doi:
471 10.1126/science.1167742.
- 472 B. G. Link, J. C. Phelan, M. Bresnahan, A. Stueve, and B. A. Pescosolido. Public concep-
473 tions of mental illness: labels, causes, dangerousness, and social distance. *Am J Public*
474 *Health*, 89(9):1328–1333, 1999. doi: 10.2105/AJPH.89.9.1328.
- 475 K. McKelvey and F. Menczer. Design and prototyping of a social media observatory. In
476 *Proc. 22nd Intl. Conf. on World Wide Web (WWW) Companion*, pages 1351–1358, 2013a.
- 477 K. McKelvey and F. Menczer. Truthy: Enabling the Study of Online Social Networks. In
478 *Proc. 16th ACM Conference on Computer Supported Cooperative Work and Social Computing*
479 *Companion (CSCW)*, 2013b.

- 480 E. F. Moran, S. L. Hofferth, C. C. Eckel, D. Hamilton, B. Entwisle, J. L. Aber, H. E.
481 Brady, D. Conley, S. L. Cutter, K. Hubacek, et al. Opinion: Building a 21st-century
482 infrastructure for the social sciences. *Proceedings of the National Academy of Sciences*,
483 111(45):15855–15856, 2014.
- 484 F. Morstatter, J. Pfeffer, H. Liu, and K. M. Carley. Is the Sample Good Enough? Com-
485 paring Data from Twitter’s Streaming API with Twitter’s Firehose. In *Proc. 7th Intl.*
486 *AAAI Conf. on Weblogs and Social Media (ICWSM)*, 2013.
- 487 A. Nematzadeh, E. Ferrara, A. Flammini, and Y.-Y. Ahn. Optimal network modularity
488 for information diffusion. *Phys. Rev. Lett.*, 113:088701, 2014. doi: 10.1103/PhysRevLett.
489 113.088701.
- 490 Open Science Collaboration. Estimating the reproducibility of psychological science.
491 *Science*, 349(6251), 2015. doi: 10.1126/science.aac4716.
- 492 J. Qiu, S. Jha, A. Luckow, and G. C. Fox. Towards hpc-abds: An initial high-performance
493 big data stack. In *Proceedings of 1st ACM Big Data Interoperability Framework Workshop:*
494 *Building Robust Big Data ecosystem*, 2014.
- 495 M. Raento, A. Oulasvirta, and N. Eagle. Smartphones: An emerging tool for so-
496 cial scientists. *Sociological Methods & Research*, 37(3):426–454, 2009. doi: 10.1177/
497 0049124108330005.
- 498 S. Rafaeli. Interactivity: From new media to communication. *Sage annual review of*
499 *communication research: Advancing communication science*, 16(CA):110–134, 1988.
- 500 J. Ratkiewicz, M. Conover, M. Meiss, B. Gonçalves, A. Flammini, and F. Menczer. De-
501 tecting and tracking political abuse in social media. In *Proc. 5th Intl. AAAI Conf. on*
502 *Weblogs and Social Media (ICWSM)*, 2011a.
- 503 J. Ratkiewicz, M. Conover, M. Meiss, B. Gonçalves, S. Patil, A. Flammini, and F. Menczer.
504 Truthy: Mapping the spread of astroturf in microblog streams. In *Proc. 20th Intl. World*
505 *Wide Web Conf. Companion (WWW)*, 2011b.
- 506 D. M. Romero, B. Meeder, and J. Kleinberg. Differences in the mechanics of information
507 diffusion across topics: Idioms, political hashtags, and complex contagion on Twitter.
508 In *Proc. 20th Intl. Conf. on World Wide Web (WWW)*, pages 695–704, 2011. doi: 10.1145/
509 1963405.1963503.
- 510 D. Ruths and J. Pfeffer. Social media for large studies of behavior. *Science*, 346(6213):
511 1063–1064, 2014. doi: 10.1126/science.346.6213.1063.
- 512 S. Sanfilippo. Redis, 2016. URL <http://redis.io/>. Last accessed April 05, 2016.
- 513 D. Selassie, B. Heller, and J. Heer. Divided edge bundling for directional network data.
514 *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2354–2363, 2011. doi:
515 10.1109/TVCG.2011.190.
- 516 M. Á. Serrano, M. Boguná, and A. Vespignani. Extracting the multiscale backbone of
517 complex weighted networks. *Proceedings of the National Academy of Sciences*, 106(16):
518 6483–6488, 2009.

- 519 A. Solem and Contributors. Celery, 2016. URL <http://www.celeryproject.org/>. Last
520 accessed April 05, 2016.
- 521 V. Subrahmanian, A. Azaria, S. Durst, V. Kagan, A. Galstyan, K. Lerman, L. Zhu, E. Fer-
522 rara, A. Flammini, F. Menczer, A. Stevens, A. Dekhtyar, S. Gao, T. Hogg, F. Kooti,
523 Y. Liu, O. Varol, P. Shiralkar, V. Vydiswaran, Q. Mei, and T. Hwang. The DARPA
524 Twitter bot challenge. *IEEE Computer*, 49(6):38–46, 2016. doi: 10.1109/MC.2016.183.
- 525 O. Tene and J. Polonetsky. Privacy in the age of big data: a time for big decisions.
526 *Stanford Law Review Online*, 64:63, 2012.
- 527 P. Terna. Simulation tools for social scientists: Building agent based models with swarm.
528 *Journal of artificial societies and social simulation*, 1(2):1–12, 1998.
- 529 The Apache Software Foundation. Apache HBase, 2016a. URL [http://hbase.apache.](http://hbase.apache.org/)
530 [org/](http://hbase.apache.org/). Last accessed April 05, 2016.
- 531 The Apache Software Foundation. Hadoop, 2016b. URL <http://hadoop.apache.org/>.
532 Last accessed April 05, 2016.
- 533 The Apache Software Foundation. Apache Solr, 2016c. URL [http://lucene.apache.](http://lucene.apache.org/solr/)
534 [org/solr/](http://lucene.apache.org/solr/). Last accessed April 05, 2016.
- 535 The Apache Software Foundation. Apache Hadoop YARN, 2016d. URL [http://hadoop.](http://hadoop.apache.org/docs/current/hadoop-yarn/hadoop-yarn-site/YARN.html)
536 [apache.org/docs/current/hadoop-yarn/hadoop-yarn-site/YARN.html](http://hadoop.apache.org/docs/current/hadoop-yarn/hadoop-yarn-site/YARN.html). Last ac-
537 cessed April 05, 2016.
- 538 J. Travers and S. Milgram. An experimental study of the small world problem. *Sociome-*
539 *try*, pages 425–443, 1969.
- 540 Twitter, Inc. Developer policy. Available at: [https://dev.twitter.com/overview/](https://dev.twitter.com/overview/terms/policy)
541 [terms/policy](https://dev.twitter.com/overview/terms/policy), Internet Archive: [https://web.archive.org/web/20160311122344/](https://web.archive.org/web/20160311122344/https://dev.twitter.com/overview/terms/policy)
542 <https://dev.twitter.com/overview/terms/policy>, 2016. Last accessed:
543 04/09/2016.
- 544 O. Varol, E. Ferrara, C. Ogan, F. Menczer, and A. Flammini. Evolution of online user
545 behavior during a social upheaval. In *Proc. ACM Web Science Conference (WebSci)*, 2014.
- 546 E. Vayena, M. Salathé, L. C. Madoff, and J. S. Brownstein. Ethical challenges of big data
547 in public health. *PLoS Comput Biol*, 11(2):e1003904, 2015. doi: 10.1371/journal.pcbi.
548 1003904.
- 549 L. Weng and F. Menczer. Topicality and impact in social media: Diverse messages,
550 focused messengers. *PLoS ONE*, 10(2):e0118410, 2015.
- 551 L. Weng, A. Flammini, A. Vespignani, and F. Menczer. Competition among memes in a
552 world with limited attention. *Sci. Rep.*, 2(335), 2012.
- 553 L. Weng, F. Menczer, and Y.-Y. Ahn. Virality prediction and community structure in
554 social networks. *Sci. Rep.*, 3(2522), 2013a.

- 555 L. Weng, J. Ratkiewicz, N. Perra, B. Gonçalves, C. Castillo, F. Bonchi, R. Schifanella,
556 F. Menczer, and A. Flammini. The role of information diffusion in the evolution of
557 social networks. In *Proc. 19th ACM SIGKDD Conference on Knowledge Discovery and*
558 *Data Mining (KDD)*, 2013b.
- 559 L. Weng, F. Menczer, and Y.-Y. Ahn. Predicting successful memes using network and
560 community structure. In *Proc. Eighth International AAAI Conference on Weblogs and*
561 *Social Media (ICWSM)*, 2014.
- 562 T. B. Wiggins, X. Gao, and J. Qiu. IndexedHBase, 2016. URL [http://salsaproj.](http://salsaproj.indiana.edu/IndexedHBase)
563 [indiana.edu/IndexedHBase](http://salsaproj.indiana.edu/IndexedHBase). Last accessed April 05, 2016.
- 564 T.-L. Wu, B. Zhang, C. A. Davis, E. Ferrara, A. Flammini, F. Menczer, and J. Qiu. Scalable
565 query and analysis for social networks: An integrated high-level dataflow system with
566 pig and harp. In M. Thai, H. Xiong, and W. Wu, editors, *Big Data in Complex and Social*
567 *Networks*. Chapman and Hall/CRC, 2016. Forthcoming.
- 568 M. Zimmer. The Twitter archive at the library of congress: Challenges for information
569 practice and information policy. *First Monday*, 20(7), 2015.