#swineflu: The Use of Twitter as an Early Warning and Risk Communication Tool in the 2009 Swine Flu Pandemic

PATTY KOSTKOVA, University College London

MARTIN SZOMSZOR, Digital Science, London

CONNIE ST. LOUIS, City University, London

The need to improve population monitoring and enhance surveillance of infectious diseases has never been more pressing. Factors such as air travel act as a catalyst in the spread of new and existing viruses. The unprecedented user-generated activity on social networks over the last few years has created real-time streams of personal data that provide an invaluable tool for monitoring and sampling large populations. Epidemic intelligence relies on constant monitoring of online media sources for early warning, detection, and rapid response; however, the real-time information available in social networks provides a new paradigm for the early warning function.

The communication of risk in any public health emergency is a complex task for governments and health-care agencies. This task is made more challenging in the current situation when the public has access to a wide range of online resources, ranging from traditional news channels to information posted on blogs and social networks. Twitter's strength is its two-way communication nature — both as an information source but also as a central hub for publishing, disseminating and discovering online media.

This study addresses these two challenges by investigating the role of Twitter during the 2009 swine flu pandemic by analysing data collected from the SN, and by Twitter using the opposite way for dissemination information through the network. First, we demonstrate the role of the social network for early warning by detecting an upcoming spike in an epidemic before the official surveillance systems by up to two weeks in the U.K. and up to two to three weeks in the U.S. Second, we illustrate how online resources are propagated through Twitter at the time of the WHO's declaration of the swine flu "pandemic". Our findings indicate that Twitter does favour reputable t bogus information can still leak into the network.

Categories and Subject Descriptors: J.3 [Health]

General Terms: Experimentation

Additional Key Words and Phrases: Epidemic intelligence, data mining, social media, global health and well-being, monitoring spread of disease, swine flu 2009 analysis, Twitter, real-time data management and public health response

ACM Reference Format:

Kostkova, P., Szomszor, M., and St. Louis, C. 2014. #swineflu: The use of twitter as an early warning and risk communication tool in the 2009 swine flu pandemic. *ACM Trans. Manage. Inf. Syst.* 5, 2, Article 8 (July 2014), 25 pages.

DOI: <http://dx.doi.org/10.1145/2597892>

1. INTRODUCTION

Communicable diseases in the 21st century remain a significant public health threat due to a number of factors: globalisation, the emergence of new diseases, and the reappearance of older infectious disease. There are many reasons for this, including frequent international air travel, growing resistance to antibiotics, more human-to-animal contact, as well as the threat of bio-terrorism. As a result, new public health measures are urgently needed. The SARS outbreak in 2003 illustrated how quickly a

Authors' addresses: P. Kostkova (corresponding author), University College London, London, UK; email: p.kostkova@uci.ac.uk; M. Szomszor, Digital Science, London, UK; C. St. Louis, City University, London, UK. Permission to make digital or hard copies of part or all 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. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author.

2014 Copyright is held by the Owner/Author. 2158-656X/2014/07-ART8 \$15.00

DOI: <http://dx.doi.org/10.1145/2597892>

new virus could globally spread. The response to these infectious disease threats needs to be rapid, evidence based, and internationally coordinated. This in turn requires instant information dissemination to enable a rapid response to public health threats at national and international levels. Information technology now provides the tools to support large-scale population monitoring. Early warning and response systems, along with rapid communication, ensure that knowledge and scientific expertise is mobilised to protect citizens. However, the unprecedented growth of user-generated information on the Internet and recently on social networks such as Facebook and Twitter has revolutionised information creation, dissemination, sharing, and management for public health needs.

Social media, such as blogging, social networking, and wikis, has attracted a great deal of interest recently as a possible source of data for Epidemic Intelligence (EI). The real-time nature of microblogging and status updates presents a unique opportunity to gather information on large numbers of individuals as well as offering the opportunity to enhance early warning outbreak detection systems. While traditional EI systems such as GPHIN and Medisys are well established and used routinely by the European Centre for Disease Control (ECDC) and the World Health Organization (WHO), new sources of data are constantly under review. Recent work [Google 2014] by companies such as Google has demonstrated that online search queries for keywords relating to flu and its symptoms can serve as a proxy for the number of individuals who are sick. The potential of Google Flu Trend for tracking flu by comparing the signal to U.S. outpatient ILI surveillance network (ILINet) was documented by Cook et al. [2011]. However, this search data remains proprietary and is therefore not available for research or for the construction of noncommercial applications. Twitter data, of a certain volume (currently 1% of the “firehouse” RT data stream is available free of charge) is publicly available and offers access into people’s online and offline real-time activity. Exploring the role of Twitter for real-time large-scale population monitoring, in order to predict the spread of diseases and provide an early warning to public health authorities, was the first aim of our research study.

The second challenge is health communication. In the past, during emergencies and public health-risk situations, news organizations and governments had complete control over what was published by the media and how it was accessed. With the relatively recent invention of the Web and increasingly popular social media, any individual can create and post material online, potentially connecting with the global population without editorial comment or moderation. While editors or, in the case of science, a peer-review process, moderate traditional news sources and scientific outlets, online publishing such as blogs, podcasting, and vlogging (video blogging¹) has enabled unverified sources of information to be published. The explosion in popularity of social media has subsequently raised concerns about the quality of information that is present online. Consumers of information on the Web now have the difficult task of assessing the quality of the information they see without any specific training or guidance. In particular, health reports online have become a hotbed for fearmongering and false advertising. These unscrupulous sites prey on the fears of the public and often exploit them for financial gain. A recent example is the widespread support that can be found online which ignores scientific evidence and claims that there is a link between MMR (measles, mumps, and rubella vaccine) and autism. Its impact can be seen in the U.K. by the increased cases of mumps and whooping cough and the return of previously eradicated infectious diseases, such as measles.

Thus, in this article, we investigate Twitter by analysing data collected from the SN (to illustrate the power of the real-time nature of the SN for early warning and also

¹http://en.wikipedia.org/wiki/Video_blogging

using the SN as a conduit to instantly disseminate news through the network). Our analysis of Twitter data from May to December 2009 demonstrates its potential as a information source for early warning systems. The study of the dissemination of the WHO's change of status of swine flu alert from "epidemic" to "pandemic" in June 2009 through the social networks also indicates that coverage is an important factor for risk communication.

The article is organized as follows: Section 2 gives the background to the study, Section 3 discusses the datasets used in the study and also contains a basic analysis of the tweets and users. Section 4 contains the first major contribution: the demonstration of the role of Twitter as an early warning system during the 2009 swine flu pandemic, while Section 5 discusses the second part of our study that considers the role of Twitter in the dissemination of information of public health importance and the quality of the media coverage. Section 6 is dedicated to discussion and future work, while Section 7 concludes.

2. BACKGROUND - GLOBAL DISEASE SURVEILLANCE AND EPIDEMIC INTELLIGENCE

In this section we discuss the background to infectious disease surveillance and epidemic intelligence. We begin by examining the role of Epidemic Intelligence (EI) and social media for influenza surveillance and continue by investigating the role of social networks in information dissemination.

Epidemic Intelligence (EI) has been defined as the automated early identification of health threats and disease outbreaks, their verification, risk assessment, and investigation to inform health authorities about the required measures to protect citizens [Kaiser et al. 2006; Kaiser and Coulombier 2006; Paquet et al. 2006]. A comprehensive summary of bio-terrorist attack detection tools was carried out by Buckeridge et al. [2005]. These electronic EI systems complement traditional sentinel surveillance systems, however, with large-scale blogging, social networks, and Web 2.0. An outbreak is often discovered earlier through EI tools than the health authorities of the country concerned might even discover through traditional reporting channels. In addition, citizens' participatory surveillance is providing vital self-reported public health information and is changing the vista by providing real-time information sources. The use of EI is on rise in the U.S. [Centers for Disease Control and Prevention 2012], as well as in Europe via the multinational Influenzanet and in the U.K. portals Flusurvey [2014], and Sickweather [2014] platform.

Our focus in this research is on the changing surveillance potential for influenza. Traditionally, surveillance of influenza and influenza like illness (ILI) follows the same principles where the reporting is symptomatic. (This methodology was different in the initial phase of swine flu study in 2009, where laboratory results of samples were required.) Also, with the high prevalence of influenza (especially in winter months), this has become ideal for active citizens' participatory surveillance, leading to the development of several successful projects such as the aforementioned Flusurvey, Influenzanet, and also Flu Near You².

The next sections discuss the development of Twitter and social media in ILI surveillance (Section 2.1) and the traditional aspects of risk communication and the potential of Twitter to improve contact with citizens during emergencies (Section 2.2).

2.1. The EI Data Source and the Role of Twitter for EI and ILI Surveillance

Epidemic intelligence has historically relied on automated news media searching systems for event-based early warnings, however, this is now radically changing.

²<https://flunearyou.org/>

The use of digital epidemiology that harvests digital data sources for public health purposes brings great potential and new challenges [Salathe et al. 2012] while creating new possibilities for the use of Big Data [Hay et al. 2013]. Barboza et al. [2013] compared operational early warning systems by scanning media news and email warning systems (i.e., BioCaster, Argus, GPHIN, HealthMap, Medisys, ProMED-mail, Puls systems discussed shortly) on the detection of A/H5N1 influenza events. They highlighted the need for “more efficient synergies and cross-fertilization of knowledge and information”. In addition, the roadmap for digital disease surveillance that incorporated new data sources was also recently outlined [Kostkova 2013], identifying six types of data sources for EI:

- (1) news/online media,
- (2) digital traces,
- (3) ProMED,
- (4) labs/clinical reports,
- (5) participatory systems, and
- (6) social media.

Although our study investigates the role of social media, namely Twitter, in early warning systems, it is important to add context with a brief discussion of other event-based surveillance data sources.

- News and online media.* These have been used as sources in EI for over a decade. Traditional systems such as Global Public Health Intelligence Network (GPHIN) [Open Text 2014], developed by Health Canada and used by WHO, as well as Medisys³, developed by the JRC, gather news from the global media to identify disease outbreak threats using multilingual natural language processing and an appropriately weighted set of keywords, categories, and taxonomies [WHO 2014; Linge et al. 2009]. Several comparative studies have looked at event-based monitoring systems. They involved examining unstructured event-based reports from GPHIN [Open Text 2014], HealthMap [2014], and EpiSPIDER [2014], then analysing them for effective global infection disease surveillance and development. Keller et al. [2009] found that, while news monitoring provides a robust EI data source, it is not suitable for very early warning systems as it usually takes several days for an event to become news and therefore news monitoring might not perform well in local disease outbreak coverage because this may never be important for mainstream media. Further, countries with state-controlled news provide unreliable information for EI; and as a result, this important role is filled by social media.
- Digital traces.* Digital traces are increasingly becoming essential signal sources that add value by providing additional information. They search keywords, loyalty purchase cards, sensor networks, drugs purchases, and mobile phone data. However, these systems typically rely on private company datasets that are not easily accessible for research. Google’s Flu Trends research has estimated an upcoming flu epidemic faster than CDC surveillance data that evaluates online search queries for keywords relating to flu [Google 2014]. According to Cook et al. [2011], this proprietary work by Google Flu Trends provides no means for verification or direct comparison. Ginsberg et al. [2009] illustrated an automated method for defining ILI-related keywords without prior knowledge of influenza. A similar study done on a smaller scale using infection experts to advise on keyword selection and correlations was conducted by Wiseman et al. [2010] on the NeLI/NRIC portal, identi-

³<http://medusa.jrc.it/medisys/homeedition/en/home.html>

fying user information needs during the swine flu pandemics in 2009 from Weblog searches. However, online search systems don't typically reflect the real-time information needs of citizens. In the field of mobile application technology, EpiCollect provides an example of mobile phone data usage for epidemiology [EpiCollect 2014; Aamesen et al. 2009].

The third type of data source is the email-based system ProMED-mail that has been a long-established informal source of emergencies utilised by infectious disease professionals. Technically, it can't be categorized under the category of "participatory surveillance systems", as it is a mailserv. However, ProMED has an element of active sharing and participation [ProMED 2010]. Additionally, as a human-moderated data source, it is also subject to bias and has a lower global coverage.

- *Labs and clinical reports.* These have historically been the backbone of surveillance systems. Microbiology laboratories contribute to surveillance by providing data of the highest reliability by means of microbiological confirmation of unusual disease patterns and specimens, albeit at the expense of timelines. However, these results remain in the domain of the government and are normally unshared with stakeholders and researchers outside the public health agencies, and, because of the long timeframes involved, are not suitable for early warning systems and new diseases surveillance.

In this study, we investigated the impact of the increased amount of Web 2.0 and user-generated content via social networking tools such as Facebook and Twitter, providing EI systems with a highly accessible source of real-time online activity. While participatory systems and social media partially overlap in terms of functionality, for the needs of this research we understand participatory systems as dedicated surveillance applications (web-based or mobile) requiring proactive participation in regularly sharing disease symptoms collected in a data-structured format. Examples are the aforementioned multilingual EpiWorks project Influenzanet [2014] and Sickweather [2014]). While participatory systems provide surveillance with valuable longitudinal datasets, they rely on voluntary participation and engagement of citizens in the application and — more importantly — their long-term retention. Also, unlike social media sharing, participatory systems typically limit submissions to a set of symptoms, thus making a trade-off between reducing coverage and the ease of contribution.

- *Social media.* Finally, our focus is social media — a Big Data source that revolutionised the speed and timeliness of EI. Facebook's privacy setting allows users to restrict their profile content and activity. However, Twitter [2014] is available in the public domain and therefore freely searchable and analyzable using a provided API [Williams 2014]. The information posted on Twitter describes a real-time activity due to the social nature of the service, unlike search queries collected by search engines. Therefore, utilizing this increasingly popular freely available data source has a potential for EI and other rapid information intelligence systems.

Furthermore, our focus is on Twitter as it provides an excellent way to sample large populations. In terms of epidemic intelligence, Twitter can be used to both track [Lampos et al. 2010; Lampos and Cristianini 2010] and predict [Szomszor et al. 2010] the spread of infectious diseases, as we demonstrated in our previous preliminary study. Lampos and Cristianini in their follow-up study [Lampos and Cristiani 2012] used their technique of supervised learning for "nowcasting" events by exploring geo-located Twitter signals in two case studies, namely ILI rates and rainfall. Further, a number of approaches adopted during the 2009 swine flu were discussed by BMJ with input from public health agencies' EI experts, highlighting the potential and practical

challenges [St. Louis and Zorlu 2012; Malik 2011]. Signorini et al. [2011] evaluated user sentiment during the swine flu in the U.S. and Influenza-Like Illness (ILI) reported disease levels. Unlike in our case, the data collection in this project started in October 2009, thus missing the first spike in the swine flu 2009 epidemics. While their method differs in using SVM for classification and focuses on the U.S., their results, in particular for predicting outbreaks around two weeks earlier than PH agencies, confirm ours.

Moreover, the *medi+board* project developed a generic dashboard for public health professionals integrating multiple data sources to enhance early-warning risk assessment and epidemic intelligence [Kostkova et al. 2014]. Further, a pandemic toolkit and simulator for state surveillance were developed in Texas but seemed to lack scientifically published results [University of Texas 2013]. ILIs were tracked and correlated with CDC surveillance data also by Culotta [2010] and a dengue fever was tracked using Twitter in Brazil by Gomide et al. [2011]. Culotta's approach used regression (while we adopted normalized cross-correlation) and illustrated strong correlation of the two datasets. The role of travel for seasonal transmission of A(H1N1) was also investigated by Balcan et al. [2009] to provide evidence for potential travel restrictions for policy makers. Recently, Salathe et al. [2013] illustrated the role of digital epidemiology and Twitter for understanding the new strain of Influenza A (H7N9) and the coronavirus (MERS-CoV).

Other than in public health domains, Twitter has also proved to have excellent real-time benefits. Earthquake detection [Sakaki et al. 2010] is made possible by examining the tweets of users in the local area containing terms related to earthquakes. When natural disasters strike, Twitter can help coordinate rapid responses [Vieweg et al. 2010] and increase situational awareness with users providing important information on local conditions (such as weather, visibility, road conditions, etc.). Both Facebook and Twitter are becoming increasingly more popular for raising awareness and raising funds for global relief [FastCompany 2014].

2.2. The Role of Social Networks for Risk Communication and News Dissemination

Twitter and SN can indeed communicate in two ways: they perform very well as a hub and disseminator and have proven to play an increasingly important role for risk communication and media coverage dissemination.

Traditionally, risk communication was conducted using mainstream media, namely TV, press, and radio. Swineflu 2009 was a breakthrough in this regard, as illustrated by Duncan [2009] in his swine flu media coverage in the EU while a specific study to analyze risk perception and information seeking behaviour during the 2009 pandemic was conducted in Germany [Walter et al. 2012]. Twitter was also investigated from the perspective of its role as a social network and news media [Kwak et al. 2010] and also as a medium for understanding concerns about public health [Paul and Dredze 2011].

Much computing research so far, such as by Kwak et al. [2010], has focused on understanding how information cascades through the Twitter network. Since Twitter users usually follow other Twitter users to stay up-to-date with what people in the social network are doing, the connections that people make are intrinsic to the dynamics of information flow. Various studies [Bakshy et al. 2011; Cha et al. 2010; Lee et al. 2010; Lerman and Ghosh 2010] show that influential people in the Twitter network (i.e., those with large numbers of followers) are the main hubs and control the spread of information. However, having large numbers of followers does not guarantee that information will propagate through the network [Cha et al. 2010]. Instead, other factors such as timeliness, accuracy, and entertainment play an import role. Elsewhere, research has centered on making sense of the information appearing in Twitter [Savage 2014]. In [Laniado and Mika 2010] Laniado and Mika analyze the use of

hashtags (terms prefixed with the “#” symbol), uncovering a complex picture of an emergent, decentralized system where semantic identifiers serve as important information markers. Interesting situations often arise because of the message limit placed on tweets. For example, new natural language processing techniques [Sriram et al. 2010] are required to cater for short messages and the small amount of information they contain.

Based on these results, applications in the medical domain started emerging. A recent study of antibiotic understanding on Twitter [Scanfeld et al. 2010] showed that social media is a useful way to disseminate medical information, but that it is also prone to abuse. In particular, Twitter can be used to assess public knowledge (e.g., the widely held but incorrect assumption that antibiotics will treat a cold, or that a course of antibiotics can be stopped once the symptoms have disappeared) and therefore reveal gaps in public understanding. Also, the problem of inaccuracy of self-diagnosis of influenza due to media hype was illustrated on the swine flu 2009 pandemic by Jutel et al. [2011]. Further, attention has turned to Web 2.0 tools other than Twitter. For example, Wikipedia could also be used for public health promotion [Heilman et al. 2011], while data correlation of news searches made by both the public and professionals was investigated in the NeLI/NRIC datasets [Kostkova 2013]. Despite these results, analysis of public health information reported on social media due to the fast nature of information during emergencies remains challenging and our research investigates this gap.

3. TWITTER DATASET AND ANALYSIS

Twitter users are free to follow any other users but do not have to be followed back (this relationship is asymmetrical) and to use this facility to build networks that support social, business, and academic activities. In June 2012, Twitter reached 500 million users [TechCrunch 2014].

When evaluating high-noise large-tweet databases, it is important to understand how “useful” the collected dataset was and what were the core sets used for analysis. In our study, the subset of over 25,000 tweets out of 3 million called “self-diagnosed” was used in the early warning part of the study, whereas for the second part of the study, tweets sent on 11/6/09 were evaluated and 31% (over 30,000) were found containing the term “pandemic” (Figure 3).

3.1. Dataset Specification

Using the standard Twitter API, we searched for the term “flu” and collected over 3 million tweets in the period from May 7th 2009 until December 22nd 2009 and carried on collecting them on a one-minute basis. We found just less than 3 million tweets containing the keyword “flu”, including individuals reporting flu symptoms or self-diagnosing, sharing links to news articles, Web sites, and blogs, and generally commenting on the topic. Enhancing our previous pilot analysis [de Quincey and Kostkova 2010], we show the most popular words in these tweets and their frequencies in Table I.

3.2. Classification of Tweets

In order to investigate the use of Twitter for EI and risk communication, we classified the tweets using the following classes (note it is possible for a tweet to be placed in more than one class).

- (1) *Tweets containing a link/url.* A popular activity in Twitter is to post a link to a Web site. Many use this mechanism to link their followers to online news articles, blogs, videos, images, etc. Because of the 140-character limit of tweets and the typically

Table I. Top 20 Most Frequently Occurring Words

Freq	Word	Freq	Word
2,993,022	flu	92,999	#swineflu
1,621,782	swine	88,801	cases
264,903	rt	82,130	#h1n1
223,876	h1n1	71,323	today
195,163	vaccine	69,071	shots
156,658	shot	66,167	hope
109,995	health	64,271	feel
107,675	sick	63,732	school
97,889	news	61,004	:(

Fig. 1. A time series showing all tweets containing the keyword “flu”, those containing links, those reporting flu, and retweets. A seven-day moving-window average has been applied to smooth the data.

long length of urls, url shortening services (such as bit.ly and tinyurl.com) are often used.

- (2) *Retweets*. Another popular Twitter behaviour is to *retweet* a message. In essence, users who see an interesting tweet will pass it onto their followers by reposting the original message and quoting the original author. Retweets themselves often contain links. We search for “rt @<hyperlink to user>” to find retweets.
- (3) *Self-Reporting flu*. We check the text of each tweet and search for phrases that indicate the user has the flu. These include the phrases “have flu”, “have the flu”, “have swine flu”, and “have the swine flu” in present and past tenses.

Figure 1 contains a time-series plot for the total number of tweets recorded during the period 11-05-2009 until 20-12-2009. A seven-day moving-window average is applied to smooth the data. The plot shows the total number of tweets containing the keyword flu (labeled “All Flu Tweets”) for each day, the total number of tweets containing a link (“Contains Link”), the total number of tweets reporting flu (“Self-Reporting via Twitter”), and the total number of retweets (“Retweets”). Due to technical problems, a section of data is missing for the period 30/08/2010 to 14/09/2010. The time series indicates significant increases in activity around week 30 (20/07/2010) and again

Fig. 2. A plot showing the proportion of links each day that contain a link and the those that are retweets.

around week 40 (28/08/2010). Posting of links constitutes the most significant percentage of tweets at around 67%, while the number of self-reporting tweets is around 5% and the number of retweets is approximately 2%.

Overall, tweets containing links are the most prominent, accounting for $\sim 65\%$ of all tweets that contain the term “flu”. Retweeting is a scarce activity, on average only 1% of tweets were retweets. Previous work [Szomszor et al. 2011] has shown there is a slight increase in retweeting activity over 2009, rising from around 1% in May 2009 to 3% by December 2009. It is not clear from the data we have gathered whether this increase in retweeting is a trend specific to flu-related tweets or a trend across the whole of Twitter. The latter seems more likely since individuals have become more aware of the retweeting practice, as well as the attributable inclusion of a retweeting functionality in the majority of Twitter clients. We also investigated tweets containing hashtags, however, these were quite rare ($\sim 10\%$) as the dataset with only one significant increase in activity was found on 16/08/09 (DOY 228).

3.3. Distribution of Links/URLs and Retweets

Since the posting of links makes up a significant proportion of flu-related tweets, we decided to perform further analysis of these cases to identify any global trends. An increase in the posting of links could indicate an increased reaction to news and other online media. Figure 2 plots the percentage of tweets for each day that contain a link (using the left axis) and the percentage of tweets that are retweets (right axis). The plot shows that the posting of links remains relatively constant over time (around 67%). The percentage of retweets displays an overall increase from approximately 0.75% in week 25 to around 3% in week 52. It is not clear from the data we have gathered whether this increase in retweeting is a trend specific to flu-related tweets or a trend across the whole of Twitter. The latter seems more likely since individuals have become more aware of the retweeting practice in Twitter since the beginning of 2009.

Fig. 3. Overall Twitter activity (left axis) versus the number of tweets containing the term “pandemic” (right axis).

3.4. Term Frequency Analysis

Further, we also performed a day-by-day term frequency analysis to determine whether specific terms exhibit a “bursty” behaviour. Much like the link analysis described before, sudden bursts in the appearance of a particular term would indicate an important event. Using a standard dictionary of stop-words and a filtering process that removes punctuation marks from each word in a tweet (i.e., !\”\$%&’()*+,-./:;<=>?@[\\]^_`{|}~), a daily total for each of the top 100 terms appearing in our sample dataset was calculated. Only one term showed a significant event, namely the term “pandemic” was used in 31% of tweets containing the term “flu” on 11/06/09 (DOY 162), as depicted in Figure 3. It is worth noting that the increased use of the term “pandemic” started one day before the status change as people and news articles speculate on the upcoming change. Further analysis of this specific event can be found later in Section 5.3.

3.5. Spam Detection Method

While we do not aim to conclude any real numbers but rather single changes and propagation patterns, detecting spam as the widely recognized problem in Twitter is essential for accuracy and minimization of bias. When particular topics start to trend (i.e., there is a sudden increase in activity), spammers leverage the popularity of a concept for illicit promotion by posting bogus links that contain the trending term or hashtag. Our analysis was often obstructed by these spam articles. Therefore, we believe it is important to find simple and cost-effective ways to automatically identify spam resources.

The 50 most frequently appearing links posted to Twitter on the 11th and 12th June 2009 were analyzed. Each resource was manually classified as spam or not spam. In many cases this was straightforward because the url had been identified as spam either by Twitter or by a url shortening service; only a small majority had to be inspected manually. For each resource, the author-post ratio was calculated by dividing the total number of distinct authors that posted the link by the total number of times the link appears in Twitter. Essentially, this ratio is a measure of how reputable the link is, with 1 being the most reputable and 0 being the least reputable. In cases of spam, the generally adopted method is for one user to repeatedly make the same post. As a result, the ratio becomes lower and lower as more posts are made.

Fig. 4. Spam detection of resources linked. Red (square) points indicate spam; green (triangles) points represent legitimate resources.

As depicted in Figure 4, after calculating the author-post ratio for the top 50 links the appeared on the 11th and 12th June 2009, it is clear that the most spammers can be identified using this simple metric. In two of the 50 cases (one legitimate post labeled spam, one spam post labeled legitimate) an error was made. In one instance, where a spam post is not identified, the spammer used multiple accounts to make postings, ensuring that the author-post ratio for the resource remained high.

To summarize, spam is a massive problem for Twitter in terms of both obstructing users and research. We have demonstrated that some simple measures can be used to accurately classify spam. Since this dataset is relatively old for an online resource, it is entirely possible that spammers have already realized this weakness and accordingly adapted their tactics.

3.6. Demographic Analysis

Having analyzed the basic characteristics of the content of the tweets, we moved on to analyzing the users who posted the tweets. From the limited profile information available, we primarily focused on finding user locations. Much research that is conducted using Twitter data relies on determining the location of Twitter users (or where tweets came from). For example, epidemic intelligence applications should be able to identify local outbreaks at an early stage. While many Twitter clients (including the version for mobile platforms) supported the inclusion of global coordinates (such as latitude and longitude) in 2009, the data was too sparse to be useful in our study. While recent changes to the Twitter data model have helped solve this problem by providing built-in constructs to properly geo-tag tweets, the data in our study period is not geo-tagged. The next best available method to find the source of tweets is to examine the location field that is filled in on a Twitter user's profile page. Since this field is free text, determining the actual geographical location (e.g., to a country level) is not straightforward. We describe our solution to this problem in the next section.

In other epidemic intelligence research [Szomszor et al. 2010], we performed text processing over the stream of tweets to discover all those indicating that the posting user had the flu. To properly evaluate the outcomes of this method, one must compare actual surveillance data reported by official health agencies (such as the Health Protection Agency in the U.K., or the Centers for Disease Prevention and Control in the U.S.) with the output from any analysis of Twitter data. To this end, a sample of 140,077

Fig. 5. Geographical breakdown of matched Twitter data.

users (all those we identified as suffering from flu during our study period) was taken and each user's profile page was crawled to determine his/her location. Hence, we use this large sample to determine the demographic breakdown of Twitter users.

We adopted the following three-stage process to convert the location string into a Geonames⁴ country code, a widely adopted standard to represent countries of the world using only two characters.

- (1) *Match the location string to a Wikipedia entry.* Google is used to search Wikipedia for the location string specified. This technique copes well with ambiguous place names (such as Cambridge, MA or Cambridge, U.K.) because users often specify context using the same term adopted in Wikipedia.
- (2) *Extract coordinates.* If a candidate Wikipedia entry has been found, the Semantic Web resource Dbpedia⁵ is queried to extract the geographical coordinates (in terms of latitude and longitude) of the entity.
- (3) *Match coordinates to Geonames code.* The coordinates extracted from the previous step are used as input to a Geonames Web service that returns the containing country's two letter code.

A total of 86,855 (or 62%) of users had specified something in the location field of their profile. Geographical coordinates were found for 45,718, of which 44,276 (33%) were recognized by the Geonames Web service. Figure 5 is a pie chart showing the geographical breakdown of Twitter users from the 44,276 that were matched. It should be noted that there is an implicit bias because only tweets that contain the term "flu" (i.e., in English only) were sampled. The U.S. accounts for the majority of traffic (60%), with the U.K. (22%) and Canada (6%) having the next biggest representation.

4. TWITTER SURVEILLANCE AND EPIDEMIC INTELLIGENCE

In this section, we will describe our algorithm investigating the self-reported tweets and present a cross-correlation with official surveillance data from the U.K. and U.S. to illustrate the potential of Twitter, due to the real-time nature of data, to act as an early warning system.

⁴<http://www.geonames.org/>

⁵<http://dbpedia.org/>

Fig. 6. Self-reporting tweets from U.K. and U.S. users.

Our research questions include the following.

- Does the usage of self-reporting tweets provide a signal indicative of the signal created by the officially reported cases?
- Can Twitter act as an early warning system and what is the measured time difference in our case study?

4.1. Self-Reporting Tweets and Correlation with U.K. and U.S. National Surveillance Data

As mentioned earlier, we first investigated the signal from Twitter users who reported having the disease (so-called self-reporting tweets). We make no estimates on how many self-reported users actually had the disease, as for EI early warning systems the signal change rather than the case numbers is what matters. Also, case numbers could be reliably confirmed only by laboratory results, thus GP symptomatic reporting could overestimate the actual disease prevalence. The geo-located Twitter self-reporting signal for the U.K. and the U.S. is depicted in Figure 6.

In order to test the accuracy of Twitter as a mechanism for self-reporting flu and hence its potential to provide early warning detection, we collected official surveillance data from the U.K. Health Protection Agency (HPA) collected by the Royal College of GPs (RCGP) [Public Health England 2013]. The HPA provides weekly reports on the RCGP influenza-like illness (ILI) consultation rate for England and Wales, Scotland, and Northern Ireland. For comparison, we calculate the percentage of tweets that are self-reporting flu for each day in our investigation period. This normalization process means that global trends in Twitter activity (e.g., spam, increased retweeting, and increased posting of links) are not factored in. Instead, the data here shows the number of individuals self-diagnosing as a percentage of all flu-related Twitter activity. The plot shown in Figure 7 contains the HPA RCGP ILI consultation rate for England and Wales (square points, right axis) and the percentage of Twitter activity reporting flu (crossed points, left axis). First impressions reveal a strong correlation between the two data sources: a sharp peak in activity on Twitter (around week 28, 6/07/2009) corresponds to the rapid increase in the number of consultations.

Fig. 7. A plot showing the RCGP ILI rate for England vs. self-reported cases on Twitter.

Fig. 8. A plot showing the U.S. ILI rate for the U.S. vs. the number of self-reported cases on Twitter.

A similar approach is to correlate U.S. users self-reporting the disease against U.S. official surveillance data taken from the CDC Web site (<http://www.cdc.gov/flu/weekly/>), the ILINet (U.S. outpatient Influenza-Like Illness surveillance Network), as illustrated in Figure 8.

4.2. Experiment: Twitter Predicts Swine Flu 2009: Normalized Cross-Correlation

The weekly time series is illustrated in Figure 7 for the U.K. and Figure 8 for the U.S., giving visual indication of the correlation and time lag, in particular, at the time of the spikes. However, in order to better quantify the correlation between two data signals (in our case, Twitter versus official surveillance data) from the “time-lag” perspective, we calculate this using the normalized cross-correlation formula.

Fig. 9. The cross-correlation plot between Twitter and the ILI reporting in the U.K. and U.S.

Therefore, to provide an indication of the correlation between Twitter and the official U.K. and U.S. surveillance data, we calculate the normalized cross-correlation ratio between various signals from Twitter and the official RCGP U.K. surveillance (collated by HPA) data from surveillance and the US from the CDC ILINet data. Since the data is gathered on a weekly basis, we perform the comparison using a weekly aggregation of Twitter data. Eq. (1) gives the normalized cross-correlation function we use, where $x(t)$ is the total number of tweets during week t , and $y(t - i)$ is the number of reported cases according to the HPA or CDC during week $(t - i)$. We calculate r across all flu tweets, that is, those that are self-reporting, those that contain links, and those that are retweets for values of i between -4 and 4 .

$$r = \frac{\sum_t (x(t) - \bar{x}) * (y(t - i) - \bar{y})}{\sqrt{\sum_t (x(t) - \bar{x})^2 * \sum_t (y(t - i) - \bar{y})^2}} \quad (1)$$

Figure 9 displays the various values of r for weekly offsets between $i = -4$ and $i = 4$. The cross-correlation ratio (or sliding dot-product) is a measure of the similarity of two signals against a moving time lag. This means that values of r for $i = 0$ represent the degree to which two signals are correlated and when $i = -1$, it represents how much the first signal predicts the second signal. Specifically, r gains values in the interval $<-1;1>$; value 0 indicates the signals are uncorrelated while value $r = 1$ indicates the signals have the same shape (although they may be of different amplitudes, which does not affect our study as estimating the number of cases is not our aim); value $r = -1$ indicates they have the same shape except they have opposite signs. The higher the value of r , the stronger the correlation. In practice, a correlation coefficient above 0.7 or 0.8 indicates a reasonably good match. Figure 9 shows that the self-reporting tweets have a strong correlation with the HPA data and CDC data. This would indicate that our filtering and normalization process has been successful, allowing us to discriminate messages that indicate someone has the flu from the general noise on Twitter. The strongest correlation occurs at $i = 0$ (when $r = 0.95$ in the U.K. and $r = 0.93$

in the U.S.), indicating a co-occurrence of tweets and surveillance data. For improving the early warning systems using social media, we need to focus on the result for negative i indicating the Twitter signal predicting the CDC/HPA data. There is still a strong correlation at $i = -1$ (when $r = 0.74$) in the U.K., indicating that the HPA surveillance data could be predicted by Twitter up to a week in advance. In the U.S. the cross-correlation is even stronger, with a prediction potential up to two weeks as we received $r = 0.87$ for $i = -1$ and still highly correlated $r = 0.75$ for $i = -2$. Therefore, this demonstrates the potential of Twitter for early warning and outbreak detection.

To summarize, having illustrated the potential of Twitter to provide an early warning up to a week before the official data from the U.K. and up to two weeks in the U.S., it is important to note that it takes about a week for the reported figures to reach the national level, be collated, and be acted upon. Therefore, the real-time monitoring of the social network could provide a warning up to two weeks earlier in the U.K. and three weeks earlier in the U.S., significantly enhancing the preparedness and response operation. However, a more robust methodology supported by disease spread models is required to identify the early warning signal reliably and avoid false positives.

4.3. Summary of Results

In this section we summarize the results to our questions.

- *Does the usage of self-reporting tweets provide a signal indicative of the signal created by the officially reported cases?* With some difference in spikes, the Twitter self-reported tweets form a function similar to the official surveillance data. As the autumn peaks in the U.K. tended to be underreported due to the introduction of the flu line (telephone service), Twitter more realistically captured this second peak.
- *Can Twitter act as an early warning system and what is the measured time difference in our case study?* Twitter can predict up to a week in the U.K. and up to two weeks in the U.S. the peak in the outbreak. As it takes around a week to report data to a national level, authorities would be informed up to two or three weeks earlier than when relying on syndromic surveillance systems.

In the next section, we will consider the role of Twitter for risk communication during a public health emergency and bring our results analyzing the role of Twitter for disseminating an urgent message, covered by online sources, during a public emergency.

5. TWITTER LINKS TO ONLINE MEDIA: RELATIONSHIP AND ANALYSIS

In this section we look in greater detail into the role of Twitter for consumption and promotion of online resources covering the swine flu 2009 outbreak and, in particular, the World Health Organization (WHO) decision to increase of the stage of the epidemic to a global “pandemic” on 11th to 12th June 2009.

This topic serves as an excellent example because it received widespread attention during 2009 and was covered extensively in the press and social media. After the WHO declaration of pandemic state, a huge volume of information was published by online media with much focus on the effectiveness of vaccination programs and the possible methods to curb the spread of infection, as shown in our preliminary study [Szomszor et al. 2011].

In particular, we seek to answer the following questions about the role Twitter plays.

- Does Twitter provide an insight into the popularity and consumption of online resources (both official and grassroots)?
- Do Twitter users have a preference when promoting online materials, for example, for official government health bodies over those of untrusted blogs?

- What are the dynamics of information dissemination during important global events? In particular, how does timeliness ultimately affect the popularity of on-line content?

5.1. Classification of Linked Resources

As highlighted earlier in Section 3.2, a significant portion of the Twitter traffic we sampled contained a link. Twitter users post links to a variety of online resources, such as news articles, blogs, videos, etc., usually because they have some interest in them and/or they want to advertise them to their followers. Therefore, analysis of the links posted on Twitter provides some insight into the interests of the Twitter population. Since the sample we have collected is focused on a particular topic (i.e., “flu”), the links posted provide a good indication of what resources are considered important by the community.

In order to investigate whether Twitter favours the dissemination of trusted information sources over untrusted ones, we conduct a classification of the most popular Web resources found in our sample dataset to find out what types of resource are the most popular.

A complete index of all hyperlinks posted to Twitter was constructed, including the total number of times the url appears as well as the total number of distinct authors. Tracking the total number of distinct authors allows us to easily distinguish spammers, as discussed in Section 3.5, and to factor out excessive self-promotion (when a user repeats the same tweet). Because of the 140-character tweet limit, many use url shortening services (such as bit.ly) to obtain a shortened version of the urls to which they wish to link. Since there are many services available to accomplish this task, a large number of different urls can point to the same resource. Hence, any url found was retrieved programmatically (using the cURL⁶ tool) to determine whether the url posted is the final destination, or if a redirection exists.

After creating an index of all resources linked, a classification task was conducted (by an experienced journalism grad student) on the most popular 769 resources posted between 02/06/09 and 29/08/09, placing each item in one of the following categories: Blog, News, Medical Organisation, Spam, Video, Poll, Comic, Aggregator, Game, or Sales, Download, Campaign, or Suspended Account.

Table II contains the total number of distinct authors and total number of resources for each classification category. The most widely represented in terms of number of distinct resources linked is spam (40%). In the majority of cases, this was simple to verify because the user’s Twitter account had been suspended, or the redirection link registered with url shortening services had been disabled.

In terms of the number of distinct authors that tweeted a reference to a resource (and hence a direct measure of its popularity), blogs are the most widely linked (26%), closely followed by official news articles (21%) and pages from official medical organizations (15%). Since blogs represent a possible source of untrusted information, we more closely analyze these.

Table III contains the top 10 most popular (in terms of the number of distinct authors that posted the link) blog resources found in our sample dataset. The most popular is a satirical piece by the popular parody newspaper *The Onion*. Other popular resources are technology related (such as Mashable and TechCrunch). However, one story that contains information contrary to the current scientific consensus did receive attention from 138 users. The article “Do NOT Let Your Child Get Flu Vaccine” is representative of the type of article that official health agencies don’t want published online, since it

⁶<http://curl.haxx.se/>

Table II. Categories of Flu-Related Resources Posted to Twitter from 02/06/09 to 29/08/09

Category	Total Authors	Total Resources
Blog	7573	162
News	6151	117
Medical Organisation	4388	38
Spam	4231	312
Video	3897	72
Poll	741	5
Comic	484	8
Aggregator	318	10
Game	294	4
Sales	288	31
Download	248	8
Campaign	63	1
Suspended account	5	1

Table III. The Most Popular Flu-Related Blog Articles Posted on Twitter from 02/06/09 to 29/08/09

URL	Total Authors
http://www.theonion.com/articles/obamas-declaration-of-swine-flu-emergency-prompts,6952/	547
http://www.benckenstein.com/digital-media/swine-flu-susan-boyle-and-the-network-multiplier-effect/	468
http://mashable.com/2009/11/10/google-flu-shot-map/	319
http://mashable.com/2009/11/14/swine-flu-appointments/	262
http://www.theatlantic.com/magazine/archive/2009/11/does-the-vaccine-matter/7723/	185
http://techcrunch.com/2009/10/26/harvard-medical-school-launches-swine-flu-iphone-app/	180
http://www.fannation.com/si_blogs/grant_wahl/posts/74041-landon-donovan-has-h1n1-flu-virus	147
http://articles.mercola.com/sites/articles/archive/2009/10/06/Why-You-Should-NOT-Vaccinate-Your-Children-Against-the-Flu-This-Season.aspx	138
http://www.informationisbeautiful.net/2009/is-the-h1n1-swine-flu-vaccine-safe/	134
http://pitchfork.com/news/35776-jens-lekman-contracts-swine-flu/	121

is not evidence based, is authored by someone with no medical qualifications, and has the potential to cause a great deal of harm.

5.2. Online Media and Twitter: Time Correlation of Flu Coverage

This section investigates a fundamental question: “Does media coverage of a certain topic cause a buzz on social media or does social media discussion cause media frenzy?” This issue was particularly important to investigate for the 2009 swine flu outbreak that experienced unprecedented media interest, however, our results were surprising.

In order to answer the question as to whether Twitter discussion influences media coverage of the disease or vice versa, we plotted the total news articles that mentioned flu from Google News against tweets in our collected database in the period of June until August 2009, when the media frenzy was at its highest due to the pandemic status. While it could be assumed that disease cases preceded media coverage or that media discussion sparked public interest causing a Twitter debate, neither time window

Fig. 10. Twitter vs. news coverage for the flu in the period of June – August 2009.

proved to be the case in our experiment. On some days, media coverage for the flu was higher where on others, Twitter discussion flourished but peaks seem synchronized, as illustrated in Figure 10.

5.3. Experiment: WHO Pandemic Status Change

In this section, we illustrate our methodology investigating media coverage and Twitter dissemination of the WHO decision to increase the status of the epidemics.

When the WHO upgraded the status of H1N1 to “pandemic” (on 11th June 2009), a significant amount of reaction was captured by Twitter. As demonstrated in Figure 3, tweets on that day contained the term “pandemic”. This event and the corresponding data we have collected provide a unique opportunity to investigate how the timeliness of the responses from major news and public health organizations, as well as how the news propagated through the network over time.

All links found in our sample dataset on the 11th and 12th June 2009 were examined. Specifically, urls were programmatically harvested to determine whether they are still active (i.e., they have not been disabled because they were spam) and if they are redirected via a url shortening service. After following all redirection links, it became apparent that many popular online news Web sites have more than one url for a particular article. For example, extra arguments are often added to the url, such as the search term used by the user to reach the page or localization information. Each resource was inspected manually to determine whether it was a direct reference to the WHO announcement. Articles from the most popular news organizations (both U.K. and U.S.) were shortlisted, along with those from two official health agencies, namely the World Health Organization (WHO) and the Center for Disease Prevention and Control (CDC). Table IV lists the important information sources we study.

Figure 11 is a plot showing the popularity of links posted to Twitter (in terms of the number of distinct authors) on an hour-by-hour basis. Ultimately, the most popular resource is the BBC article, but this is not the first to make an appearance in Twitter.

Table IV. List of News Organizations and Public Health Agencies that Reported the Pandemic Status Change in June 2009

Source	Abbreviation
BBC	BBC
CNN	CNN
REUTERS	REUT
BNO News	BNO
USA Today	USAT
Washington Post	WP
MSNBC	MS
Guardian	GUAR
Financial Times	FT
World Health Organization	WHO
Centers for Disease Prevention and Control	CDC

Fig. 11. Hour-by-hour breakdown of the most popular resources posted to Twitter from the major news organizations and public health agencies when the WHO announced H1N1 was pandemic.

As Table V shows, CNN, Reuters, and USA Today were the first to arrive in Twitter — four hours before the BBC article was picked up. Both WHO and CDC also have articles that appear in Twitter (CDC arriving much sooner than WHO), but their uptake is relatively small compared to the BBC and CNN articles. A link to the Web site of the European Centre for Disease Prevention and Control (ECDC) covering the news appeared only once.

For illustration purposes, Figure 12 illustrates the most popular retweeted link to the BBC coverage of the WHO announcement featuring a video by WHO DG Dr. Chan.

5.4. Summary of Results

Here we briefly summarize the results presented in this section.

— *Does Twitter provide an insight into popularity and consumption of online resources (both official and grassroots)?* Twitter reflects a significant proportion of the on-line population. Web resources that are of particular interest are easy to spot and analyze.

Table V. Order of Appearance of Major News Articles
Covering the Pandemic Status Change of H1N1

Time Period (11/06/09) GMT	Source
1000-1100	CNN, REUT, USAT
1100-1200	
1200-1300	BNO, CDC
1300-1400	MS, FT,
1400-1500	BBC, WP
1500-1600	GUAR
1600-1700	WHO

Fig. 12. WHO announcement covered by BBC.

- *Do Twitter users have a preference when promoting online materials, for example, for official government health bodies over those of untrusted blogs?* The articles related to swine flu that became popular were often from quality news sites (such as the BBC) or official medical organizations (such as the WHO, CDC, or ECDC). However, in some cases blog posts with poor scientific merit did become popular.
- *What are the dynamics of information dissemination during important global events? In particular, how does timeliness ultimately affect the popularity of online content?* Twitter reacts extremely quickly to online media. Within the space of a few hours, most major news organizations had published on the topic and those articles were propagated through Twitter. It would seem that timeliness isn't a good predictor of overall success: between the 11th and 12th June 2009, the BBC article became the most popular even though it appeared in Twitter four hours later than those of other news agencies.

6. DISCUSSION AND FUTURE WORK

There are a number of questions and issues to discuss.

First, how representative is Twitter usage with respect to population demographics? A recent study by the Pew Research Centre investigating Twitter usage in America [Pew Research Internet Project 2010] revealed that Internet users aged 18 to 29 are significantly more likely to use Twitter than older adults. However, minority (African American and Latino) Internet users are more than twice as likely to tweet than white Internet users. Further, women (10%) are using the service more actively than men (7%). Personal information dominates the communication (72%), closely followed by

work communication (62%). Previous research by the Pew Research Centre [Fox 2009] found that 17% of users used their mobile phones to search information about their health. While there is a bias due to gaps in usage of SNs at present, with a dramatic increase in social network usage globally, demographic bias is unlikely to be an issue for Twitter research in the near future.

In order to build on this study further, we will perform a content analysis of all articles appearing in Twitter on the 11th and 12th June 2009 to assess their journalistic merit and categorizing them in terms of the number of properly used terms (e.g., the correct use of the term “pandemic”). Accurate definitions of the virus in question along with accuracy of figures used and whether sources of information are properly quoted will also be examined. This analysis will also be performed for all news articles appearing in the English-speaking press to assess similarities and differences in the quality of information. Our hypothesis is that the online printed press will be similar to online media: some sources of information can be trusted more than others.

Location awareness remains an issue; our reliance on profile variables is introducing an inevitable bias, however, the availability of GPS coordinates of tweets, providing the desirable accuracy and location awareness, is still in the far future from a global users’ perspective.

To reliably enhance EI systems to assist public health agencies in early warning and preparedness, modeling of social media user behaviour needs to complement modeling of disease prevalence (seasonal diseases like flu with annual cycles, multiannual cycles, etc.) to avoid false positives. There are attempts going in the modeling direction [Andersen et al. 2007].

Also, full integration efforts combining all social media, traditional surveillance, and social network data streams in a single easy-to-use dashboard for public health and EI professionals are necessary to practically enhance the early warning capacity and rapid response at international public health agencies.

However, the most important question from a healthcare-risk communication perspective is understanding how effectively social media can be utilized for public health communication in comparison with other mass media (e.g., radio, health channels, TV coverage, leaflets, etc.) as well as the population’s behavioural response. This will require subsequent research. Social networks have the potential to spread evidence-based information from official healthcare sources (WHO, HPA, ECDC, etc.) and, as a consequence, to quickly create high-quality media coverage. A key challenge for agencies will be to work out how to get the attention of the majority of citizens. This should help to protect the public from relying on distorted coverage in the tabloid media that can sometimes sensationalise the issue, thus resulting in the fuelling of concerns.

7. CONCLUSIONS

Epidemic intelligence, global public health surveillance, and population monitoring are heavily reliant on information from traditional surveillance, media scanning systems, and social networks. In this article we presented our results from the swine flu 2009 study demonstrating the potential of the social network to predict peaks in the outbreak up to week before the official surveillance data in the U.K. and up to two to three weeks earlier in the U.S. As these are normally available in a collated format at national levels a week after the data are reported, this study illustrated that Twitter can provide a very early warning system, in our case study up to two weeks for the U.K. and two to three weeks for the U.S., before public health authorities could ascertain this using current systems.

Further, social network propagation of media coverage of public health events is a dynamic process. Our results demonstrated that more “reputable” and “trusted” media such as BBC are more successful in communicating the risk of a pandemic through

Twitter than less reputable outlets when WHO declared swine flu 2009 as pandemic, while communications from the public health agencies themselves were largely ignored by Twitter users. This is an important lesson for health authorities planning future risk communications.

Through collaboration with public health agencies such as HPA, ECDC, and WHO, this research aims to advance global real-time public health signals monitoring and to develop integrated platforms that can assist public health experts around the globe in protecting populations from future epidemics.

ACKNOWLEDGMENTS

We should like to acknowledge Achintya Rao for his assistance classifying the resources found in Twitter. We would also like to thank Ed de Quincey for his foresight in setting up the Twitter data collection in May 2009—this work would not be possible without it.

REFERENCES

- Aamesen, D. M., Huntley, D. M., Feil, E. J., Al-Own, F., and Spratt, B. G. 2009. EpiCollect: Linking smart-phones to web applications for epidemiology, ecology and community data collection. *PLoS ONE* 4, 9.
- Andersen, E., Bock, D., and Frisen, M. 2007. Modelling influenza incidence for purpose of on-line monitoring. Res. rep. 2007:5, Statistical Research Unit, Department of Economics, Goteborg University, Sweden.
- Bakshy, E., Hofman, J. M., Mason, W. A., and Watts, D. J. 2011. Everyone's an influencer: Quantifying influence on twitter. In *Proceedings of the 4th ACM International Conference on Web Search and Data Mining (WSDM'11)*. 65-74.
- Balcan, D., Hu, H., Goncalves, B., Bajardi, P., Poletto, C., Ramasco, J. J., Paolotti, D., Perra, N., Tizzoni, M., Van Den Broeck, W., Colizza, V., and Vespignani, A. 2009. Seasonal transmission potential and activity peaks of the influenza a(h1n1): A monte carlo likelihood analysis based on human mobility. *BMC Med.* 2009, 7:45.
- Barboza, P., Vaillant, L., Mawudeku, A., Nelson, N. P., Hartley, D. M., Madoff, L. C., Linge, J. P., Collier, N., Brownstein, J. S., Yangarber, R., and Astagneu, P. 2013. Evaluation of epidemic intelligence systems integrated in the early alert in and reporting project of the detection of a/h5n1 influenza events. *PLOS One*, 8, 3.
- Buckeridge, D. L., Burkom, H., Campbell, M., Hogan, W. R., and Moore, A. W. 2005. Algorithms for rapid outbreak detection: A research synthesis. *J. Biomed. Informatics* 38, 2, 99-113.
- Centers for Disease Control and Prevention. 2012. BioSense correlate. <http://www.cdc.gov/biosense/correlate>.
- Cha, M., Haddadi, H., Benevenuto, F., and Gummadi, K. 2010. Measuring user influence in twitter: The million follower fallacy. In *Proceedings of the International AAAI Conference on Weblogs and Social Media (ICWSM'10)*.
- Cook, S., Conrad, C., Fowlkes, A. L., and Mohebbi, M. H. 2011. Assessing google flu trends performance in the United States during the 2009 influenza virus a (h1n1) pandemic. *PLoS ONE* 6, 8.
- Culotta, A. 2010. Towards detecting influenza epidemics by analyzing twitter messages. In *Proceedings of the 1st Workshop on Social Media Analytics*. 115-122.
- de Quincey, E. and Kostkova, P. 2010. Early warning and outbreak detection using social networking web-sites: The potential of twitter. In *Proceedings of the 2nd International ICST Conference on Electronic Healthcare*, P. Kostkova Ed., Lecture Notes of the Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering, vol. 27, Springer, 21-24.
- Duncan, B. 2009. How the media reported the first day of the pandemic h1n1 2009: Results of eu-wide media analysis. *Eurosurveill.* 14, 30.
- EpiCollect. 2014. <http://www.spatialepidemiology.net/epicollect/>.
- EpiSpider. 2014. <http://www.epispider.org/>.
- FastCompany. 2014. www.fastcompany.com/blog/kit-eaton/technomix/facebook-twitter-turn-charity-efforts-11.
- Flusurvey. 2014. <http://flusurvey.org.uk/>.
- Fox, S. 2009. The social life of health information. <http://www.pewinternet.org/Reports/2009/8-The-Social-Life-of-Health-Information/14-About-Us-Methodology.aspx?view=all>.
- Ginsberg, J., Mohebbi, M. H., Patel, R. S., Brammer, L., Smolinski, M. S., and Brilliant, L. 2009. Detecting influenza epidemics using search engine query data. *Nature* 457, 1012-1014.

- Gomide, J., Veloso, A., Meira, W. Jr., Almeida, V., Benevenuto, F., Ferraz, F., and Teixeira, M. 2011. Dengue surveillance based on a computational model of spatio-temporal locality of twitter. In *Proceedings of the 3rd International Web Science Conference*.
- Google. 2014. Flu trends. <http://www.google.org/flutrends/>.
- Hay, S. I., George, D. B., Moyer, C. L., and Brownstein, J. S. 2013. Big data opportunities for global infectious disease surveillance. *PLOS Med.* 10, 4.
- Health Map. 2014. <http://healthmap.org/en/?gclid=COMHjffHtiKwCFdQLfAodcH-V-Q>.
- Heilman, J. M., Kemmann, E., Bonert, M., Chatterjee, A., Ragar, B., et al. 2011. Wikipedia: A key tool for global public health promotion. *J. Med. Internet Res.* 13, 1. <http://www.jmir.org/2011/1/e14/>.
- Influenzanet. 2014. <http://www.influenzanet.eu/>.
- Jutel, A., Baker, M. G., Stanley, J., Huang, Q. S., and Bandaranayake, D. 2011. Self-diagnosis of influenza during a pandemic: A cross-sectional survey. *BMJ Open* 1, 2.
- Kaiser, R., Coulombier, D., Maldari, M., Morgan, D., and Paquet, C. 2006. What is epidemic intelligence, and how it is being improved in Europe? *Eurosurveill.* 11, 2.
- Kaiser, R. and Coulombier, D. 2006. Different approaches to gathering epidemic intelligence in Europe. *Eurosurveill.* 11, 17.
- Keller, M., Blench, M., Tolentino, H., Freifeld, C. C., Mandl, K. D., Mawudeku, A., Eusenbach, G., and Brownstein, J. S. 2009. Use of unstructured event-based reports for global infectious disease surveillance. *Emerg. Infect. Diseases J.* 15, 5.
- Kostkova, P. 2013. A roadmap to integrated digital public health surveillance: The vision and the challenges. In *Proceedings of the 22nd International Conference on World Wide Web Companion (WWW'13Companion)*. 687–694.
- Kostkova, P., Fowler, D., Wiseman, S., and Weinberg, J. R. 2013. Major infection events over 5 years: How is media coverage influencing online information needs of health care professionals and the public? *J. Med. Internet Res.* 15, 7.
- Kostkova, P., Garbin, S., Moser, J., and Pan, W. 2014. Integration and visualization public health dashboard: The medi+board pilot project. In *Proceedings of the Companion Publication of the 23rd International Conference on World Wide Web (WWWCompanion'14)*. International World Wide Web Conferences Steering Committee, 657–662. DOI:<http://dx.doi.org/10.1145/2567948.2579276>.
- Kwak, H., Changhyun, L., Hosung, P., and Moon, S. 2010. What is twitter, a social network or a news media? In *Proceedings of the 19th International Conference on World Wide Web (WWW'10)*. 591–600.
- Lamos, V., de Bie, T., and Cristianini, N. 2010. Flu detector – Tracking epidemics on twitter. In *Proceedings of the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML/PKDD'10)*. 599–602.
- Lamos, V. and Cristianini, N. 2010. Tracking the flu pandemic by monitoring the social web. In *Proceedings of the 2nd IAPR Workshop on Cognitive Information Processing (CIP'10)*. 411–416.
- Lamos, V. and Cristiani, N. 2012. Nowcasting events from the social web with statistical learning. *ACM Trans. Intell. Syst. Technol.* 3, 4.
- Laniado, D. and Mika, P. 2010. Making sense of twitter. In *Proceedings of the International Semantic Web Conference (ISWC'10)*. 470–485.
- Lee, C., Kwak, H., Park, H., and Moon, S. 2010. Finding influentials based on the temporal order of information adoption in twitter. In *Proceedings of the 19th International Conference on World Wide Web (WWW'10)*. 591–600.
- Lerman, K. and Ghosh, R. 2010. Information contagion: An empirical study of the spread of news on digg and twitter social networks. In *Proceedings of the 4th International AAAI Conference on Weblogs and Social Media (ICWSM'10)*.
- Linge, J. P., Steinberger, R., Weber, T. P., Yangarber, R., Van Der Goot, E., Al Khudhairy, D. H., and Stilianakis, N. I. 2009. Internet surveillance systems for early alerting of health threats. *Eurosurveill.* 14, 13.
- Malik, S. 2011. Which idea is likely to make the biggest impact on healthcare by 2020? *BMJ* 2011, 342.
- Open Text. 2014. Customer stories. http://www.ot.co.uk/global/customers/successstories.htm?sys_action=show&id=661.
- Paquet, C., Coulombier, D., Kaiser, R., and Ciotti, M. 2006. Epidemic intelligence: A new framework for strengthening disease surveillance in Europe. *Eurosurveill.* 11, 12.
- Paul, M. J. and Dredze, M. 2011. You and what you tweet: Analyzing twitter or public health. In *Proceedings of the 5th AAAI Conference on Weblogs and Social Media*.

- Pew Research Internet Project. 2010. <http://www.pewinternet.org/Reports/2010/Twitter-Update-2010.aspx>.
- ProMED. 2010. <http://www.promedmail.org/>.
- Public Health England. 2013. www.hpa.org.uk.
- Sakaki, T., Okazaki, M., and Matsuo, Y. 2010. Earthquake shakes twitter users: Real-time event detection by social sensors. In *Proceedings of the 19th International Conference on World Wide Web*. 851–860.
- Salathe, M., Bengtsson, L., Bodnar, T. J., Brewer, D. D., Brownstein, J. S., Buckee, C., Campbell, E. M., Cattuto, C., Khandelwal, S., Mabry, P. L., and Vespignani, A. 2012. Digital epidemiology. *PLoS Comput. Biol.* 8, 7.
- Salathe, M., Freifeld, C. C., Mekaru, S. R., Tomasulo, A. F., and Brownstein, J. S. 2013. Influenza a (h7n9) and the importance of digital epidemiology. <http://nwjm.org>.
- Savage, N. 2014. Twitter as medium and message. *Comm. ACM* 54, 3, 18–20.
- Scanfeld, D., Scanfeld, V., and Larson, E. L. 2010. Dissemination of health information through social networks: Twitter and antibiotics. *Amer. J. Infect. Control* 3, 8, 182–188.
- Sickweather. 2014. <http://www.sickweather.com/>.
- Signorini, A., Segre, A. M., and Polgreen, P. M. 2011. The use of twitter to track levels of disease activity and public health concern in the U.S. during the influenza a h1n1 pandemic. *PLOS One* 6, 5.
- Sriram, B., Fuhry, D., Demir, E., Ferhatosmanoglu, H., and Demirbas, M. 2010. Short text classification in twitter to improve information filtering. In *Proceedings of the 33rd International ACM SIGIR Conference on Research and Development in IR (SIGIR'10)*. 841–842.
- St. Louis, C. and Zorlu, G. 2012. Can twitter predict disease outbreaks? <http://www.bmj.com/content/344/bmj.e2353>
- Szomszor, M., Kostkova, P., and de Quincey, E. 2010. #Swineflu: Twitter predicts swine flu outbreak in 2009. In *Proceedings of the 3rd International Conference on Electronic Healthcare*, M. Szomszor and P. Kostkova, Eds. Lecture Notes of the Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering, vol. 69, Springer, 18–26.
- Szomszor, M., Kostkova, P., and St Louis, C. 2011. Twitter informatics: Tracking and understanding public reaction during the 2009 swine flu pandemics. In *Proceedings of the IEEE/WIC/ACM International Conferences on Web Intelligence and Intelligent Agent Technology (WI-IAT'11)*. Vol. 1. 320–323.
- TechCrunch. 2014. <http://techcrunch.com/2012/07/30/analyst-twitter-passed-500m-users-in-june-2012-140m-of-them-in-us-jakarta-biggest-tweeting-city/>.
- Twitter. 2014. <http://www.twitter.com>.
- University of Texas. 2013. Texas pandemic flu toolkit. <http://flu.tacc.utexas.edu/>.
- Vieweg, S., Hughes, A. L., Starbird, K., and Palen, L. 2010. Microblogging during two natural hazards events: What twitter may contribute to situational awareness. In *Proceedings of the 28th International Conference on Human Factors in Computing Systems (CHI'10)*. 1079–1088.
- Walter, D., Bohmer, M. M., Reiter, S., Krause, G., and Wichmann, O. 2012. Risk perception and information-seeking behaviour during the 2009–10 influenza a (h1n1) pdm09 pandemic in Germany. *Eurosurveill.* 17, 13.
- WHO. 2014. Epidemic intelligence – Systemic event detection. <http://www.who.int/csr/alertresponse/epidemicintelligence/en/index.html>.
- Williams, D. 2014. API overview. <http://apiwiki.twitter.com/API-Overview>.
- Wiseman, S., Kostkova, P., De Quincey, E., and Jawaheer, G. 2010. Providing guidance during the swine flu outbreak in 2009: An evaluation study of the national resource for infection control (nric). Poster and abstract in *Proceedings of the 14th International Conference on Infectious Diseases (ICID'10)*.

Received December 2012; revised March 2014; accepted March 2014