

HealthTweets.org: A Platform for Public Health Surveillance Using Twitter

Mark Dredze^{1,2}, Renyuan Cheng², Michael J. Paul², David Broniatowski³

¹Human Language Technology Center of Excellence, ²Department of Computer Science
Johns Hopkins University, Baltimore, MD 21211

³Department of Engineering Management and Systems Engineering
The George Washington University, Washington, District of Columbia
mdredze@cs.jhu.edu | crenyua1@jhu.edu | mpaul@cs.jhu.edu | broniatowski@gwu.edu

Introduction

In recent years, tools for mining social media and Web data have yielded tremendous advances in public health capabilities. One area that has received considerable attention is infectious disease surveillance, with a focus on influenza (Collier 2012; Culotta 2010; Eiji Aramaki and Morita 2011; Paul and Dredze 2011b). Studies have validated the use of Twitter for influenza surveillance, with mined trends closely tracking public health agency data. e.g., the U.S. Centers for Disease Control and Prevention (CDC). Trends are produced in real time, up to two weeks before CDC influenza reports.

A gap remains between these promising results and utilization by public health officials, who lack the requisite technology and resources to produce similar trends. Several commercial websites offer services, including influenza surveillance and other health trends using social media, such as MappyHealth, Crowdbreaks, and Sickweather. However, these sites are not focused on new methods research, which could yield capabilities in new areas. Other sites, such as HealthMap.org (Freifeld et al. 2008), Flu Near You (Chunara et al. 2013), and Columbia's Influenza forecasts (Shaman et al. 2013) do not focus on social media data. Despite these excellent resources, a gap remains for delivering the latest research on public health and social media.

We present HealthTweets.org, a new platform for sharing the latest research results on Twitter data with researchers and public officials. In this demo paper, we describe data collection, processing, and features of the site. The goal of this service is to transition results from research to practice.

Data Processing

We begin with a description of how trends are mined from Twitter data, a process which produced state of the art results for influenza surveillance (Broniatowski, Paul, and Dredze 2013; Lamb, Paul, and Dredze 2013). We use the Twitter Streaming API to collect two streams: a random 1% of public tweets (SAMPLE), and a keyword stream (capped at 1% of public tweets) (HEALTH). HEALTH uses 269 keywords automatically selected from prior experiments to yield a high percentage of health tweets (Paul and Dredze 2011a).

We use a statistical classifier described in (Paul and Dredze 2011a) to identify tweets about health from HEALTH. The classifier has an estimated F-1 score of .70, and evenly balances precision and recall. Approximately 30% of tweets in HEALTH are classified as health related.

Tweets classified as health are processed by a series of automated annotators, which identify health trends. Annotators can use keywords, classifiers, rules, or arbitrary combinations of other annotators, e.g., select tweets with certain keywords, then use a classifier. We have added keywords from the HHS Now Trending Challenge ¹, and the influenza infection classifier from Lamb, Paul and Dredze (2013).

Every identified health tweet and every tweet from SAMPLE are geolocated using Carmen (Dredze et al. 2013), which attempts to identify the Country, State, County and City. Health trends are then computed based on the location, and (optionally) are normalized based on the number of tweets in a given location in SAMPLE, e.g., the rate of influenza in New York City is normalized by SAMPLE tweets from New York City. This entire process is run daily to produce updated trends, processing 10 million tweets a day with a total of over 3.4 billion tweets in the past 2.5 years.

Since data collection requires access to Twitter, gaps occur due to power loss or network interruption. We fill these gaps by estimating the missing data based on adjacent days.

Site Overview

HealthTweets.org presents health trend information in several ways. The three main visualizations are temporal trends, locations and maps. We describe each of these in detail.

Trends

Trends provides users an up to date summary and comparison of current health trends across annotators and locations. For example, a public health department could check what the current influenza rate is in their area, and compare it to the rate in other areas for the past 3 months.

Trends are shown on a 2-dimensional plot, with time (x-axis) and prevalence (y-axis). Each plot line corresponds to an annotator trend (e.g., influenza keywords) and a location (e.g., New York City). The plot can span any date range that

¹<http://nowtrending.hhs.gov/>

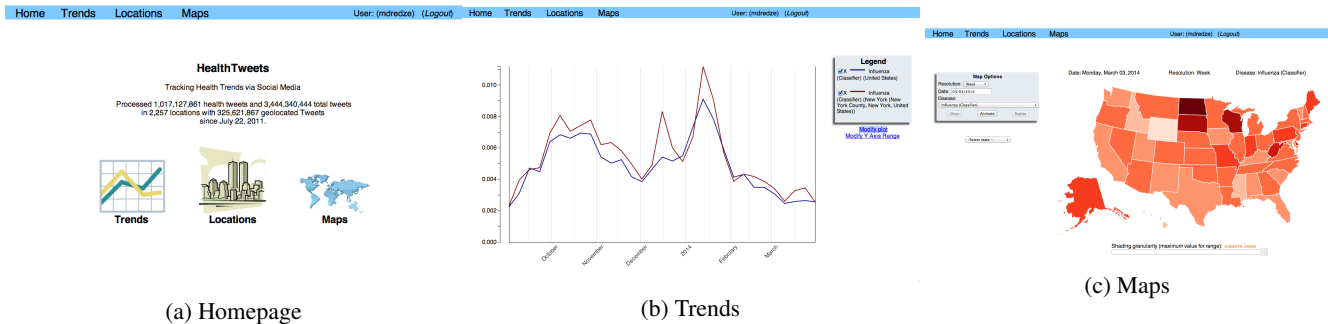


Figure 1: Screen shots demonstrating some features of the website.

is supported by the historical collection, one of three resolutions (daily, weekly or monthly aggregates) and either raw counts or normalized rates. Up to 4 plot lines can be compared at once.

Plots are dynamically created with SVG, allowing for user interaction. A user can selectively disable lines, modify the range of the y-axis, and edit the plot to add or remove lines. A permanent URL allows for sharing of constructed plots.

Locations

Locations shows detailed information about specific locations across all annotators and time. A user focused on a specific location can examine all associated health trends. For a given location and date, statistics are shown for each trend, including the rate and raw count. Statistics are summarized for the selected day, week and month, as well as the entire history in the database. Links are provided to related locations, including all known locations contained within the selected location and the encompassing geographic region.

Maps

Maps provide a geographic view of a single health trend in the United States, Europe or the World, allowing users to compare trends spatially across regions. The user selects one of the three maps, an annotator, resolution (daily, weekly, monthly), and a date. The resulting map color codes each region (country or US state) with the normalized rate of the trend. The map is interactive: when a region is selected the rate is shown as the number of health trend tweets and the total number of tweets (from SAMPLE), as well as known locations contained in the region.

In addition to generating maps for a single time period, the site creates animations of geographic trends over time. The user selects the start and end date of the animation, and the speed with which it is played. The site then generates a video, enabling for the creation of on the fly visualizations of temporal geographic trends.

Future Directions

The goal of HealthTweets.org is not to replace existing surveillance websites (described above) but to enable the sharing of state of the art methods for the exploration of and new research into social media for public health. Our goal is to attract a diverse audience of public health professionals,

who can assist in the evaluation and validation of new methods through data sharing and collaboration. As public health officials utilize data shared through the site, we hope it will aid in methodological innovations and the discovery of new capabilities for public health.

Acknowledgments M. J. Paul is supported by a Microsoft Research PhD fellowship.

References

- Broniatowski, D.; Paul, M. J.; and Dredze, M. 2013. National and local influenza surveillance through twitter: An analysis of the 2012-2013 influenza epidemic. *PLOS ONE*.
- Chunara, R.; Aman, S.; Smolinski, M.; and Brownstein, J. S. 2013. Flu near you: an online self-reported influenza surveillance system in the usa. *Online Journal of Public Health Informatics* 5(1).
- Collier, N. 2012. Uncovering text mining: A survey of current work on web-based epidemic intelligence. *Global Public Health* 7(7):731–749.
- Culotta, A. 2010. Towards detecting influenza epidemics by analyzing Twitter messages. In *ACM Workshop on Soc.Med. Analytics*.
- Dredze, M.; Paul, M. J.; Bergsma, S.; and Tran, H. 2013. Carmen: A twitter geolocation system with applications to public health. In *AAAI Workshop on Expanding the Boundaries of Health Informatics Using AI (HIAI)*.
- Eiji Aramaki, S. M., and Morita, M. 2011. Twitter catches the flu: Detecting influenza epidemics using twitter. In *Empirical Natural Language Processing Conference (EMNLP)*.
- Freifeld, C. C.; Mandl, K. D.; Reis, B. Y.; and Brownstein, J. S. 2008. Healthmap: Global infectious disease monitoring through automated classification and visualization of internet media reports. *J Am Med Inform Assoc* 15(2):150–157.
- Lamb, A.; Paul, M. J.; and Dredze, M. 2013. Separating fact from fear: Tracking flu infections on twitter. In *NAACL*.
- Paul, M. J., and Dredze, M. 2011a. A model for mining public health topics from twitter. Technical Report -, Johns Hopkins University.
- Paul, M. J., and Dredze, M. 2011b. You are what you tweet: Analyzing twitter for public health. In *ICWSM*.
- Shaman, J.; Karspeck, A.; Yang, W.; Tamerius, J.; and Lipsitch, M. 2013. Real-time influenza forecasts during the 2012–2013 season. *Nature communications* 4.