

# Monitoring Real-time Spatial Public Health Discussions in the Context of Vaccine Hesitancy

Michael C. Smith, M.S.E<sup>1</sup>, Mark Dredze, Ph.D.<sup>2</sup>, Sandra Crouse Quinn, Ph.D.<sup>3</sup>, David A. Broniatowski, Ph.D.<sup>1</sup>

<sup>1</sup>The George Washington University, Washington, DC, USA; <sup>2</sup>The Johns Hopkins University; <sup>3</sup>The University of Maryland, College Park, Maryland, USA

## Abstract

*Social media provide the potential to keep up with public discussions more quickly, at lower cost, and at potentially higher granularity and scope than do traditional surveys<sup>9</sup>. This paper details a preliminary system of real-time geographical monitoring and analysis using the context of the vaccine-hesitancy discussion across the United States, a valuable backdrop for such a system because of the diverse and impactful nature of the vaccination discussions as they appear, change, and influence the public<sup>12,20</sup>. We combine various methods in machine learning to geolocate, categorize, and classify vaccination discussions on Twitter. As a proof of concept, we show analyses with a prominent anti-vaccine discussion that validate the system with results from traditional surveys, yet also provide valuable spatial statistical power on top of such surveys on maps of the United States. We detail limitations and future work, yet still conclude that the system and the answers it enables are important because they will allow for more targeted and effective communication and reaction to the discussion as a first step towards monitoring people's views.*

## Introduction

Achieving herd immunity is a critical element of a vaccine's global effectiveness because it is vital to limiting transmission and reducing incidence. In 2007, most vaccine-preventable diseases were at an all-time low incidence<sup>28</sup>, but in recent years the spread of vaccine-preventable diseases has increased. For example, there were outbreaks of measles in Disneyland in 2014<sup>15</sup>, of pertussis in 2014<sup>34</sup>, and others. This urgent rise in such diseases has highlighted recent increases in vaccine hesitancy as prominent and controversial issues seen as potential threats to herd immunity<sup>20</sup>.

Reasons for this vaccine hesitancy, defined by the WHO as "refer[ring] to delay in acceptance or refusal of vaccines despite availability of vaccination services...complex and context specific varying across time, place and vaccines...includ[ing] factors such as complacency, convenience and confidence"<sup>16</sup>, are varied. Not only are there myriad drivers of these decisions that vary spatially<sup>12</sup>, but no "strategies intended to address vaccine hesitancy", were universally effective, if effective at all<sup>11</sup>. Because of the potential impact of these decisions, the public health community needs to track and understand these rationales as they appear, change, and influence the public. This paper details a preliminary system of real-time geographical monitoring and analysis in the context of the vaccine hesitancy discussion across the United States. The system and the answers it enables, with accompanying theoretical advantages of using social media and survey data together, are important because they will allow a first step towards more targeted and effective communication and reaction to widespread discussion, critical to reducing such hesitancy.

## Literature Review

This paper aims to produce an efficient, broad monitoring system for the discussions that exist in the United States about vaccine hesitancy. We detail the reasons for the system's potential for improvement over survey methods, and pertinent survey research as an option for validating the system.

Exploring new factors relevant to people's vaccination decisions across the nation would be slow and expensive using traditional survey methods – regardless of methods chosen, it would be difficult to keep up with new outbreaks of hesitancy or disease. Enter social media, "which provide unprecedented, realtime access to the attitudes, beliefs, and behaviors of people from across demographic groups"<sup>9</sup>. Social media is already being used to effectively, quickly, and cheaply track health issues such as disease incidence<sup>2,8,22,29</sup>. Not only do social media successfully track disease, but studies have shown that it may also track public opinion related to medical and health issues<sup>19,32</sup>. Vaccination decisions are opinions of sections of the public; therefore, social media can and should be used to track vaccine hesitancy<sup>9</sup>. This work aims to spatially track discussion about vaccination using Twitter messages as a starting point towards capturing and validating opinions; to our knowledge no system exists with such spatial granularity and capability in this context.

To validate the system one can look at ground truth of how people make decisions in similar contexts; Quinn and colleagues have produced a body of work with such aims<sup>13,24,25,27</sup>. Overall, they showed via surveys that such factors as public trust, demographics, risk perception, and social norms influence vaccine decision making<sup>25,27</sup>. For example, in a qualitative study, Quinn et al. showed that dimensions of public trust affect medical decisions in a study about postal workers' reactions during the 2001 anthrax attacks<sup>27</sup>. These "attitudinal and experience variables [and] demographic characteristics"<sup>13</sup> provide insight into how rationales about vaccination decisions may vary. They provide a means of validating the system, a starting point for exploring the spatial and sociodemographic variability in vaccination decisions, and an opportunity to confirm that hypotheses established in limited survey environments hold in wider contexts. For example, is there a significant pocket of people in a certain area who are hesitant to vaccinate because they do not trust the government?

The monitoring system able to broadly, cheaply, and quickly test such survey and spatial hypotheses in this context is the novel contribution of this work. Specifically, it is A) a processing system for classifying messages and their sentiment that B) integrates existing analyses for topic and location and C) provides an extensible framework for statistically testing spatial hypotheses about vaccine hesitancy given the generated messages and metadata on social media. This system involves using targeted methodologies and leverages theoretical advantages from using social media data in concert with survey data. How might such a system shed light on vaccine-hesitancy discussions across the USA, what are its limitations, and how could it be used as a first stepping stone to augment survey methods?

## Methods

Given our context of vaccination discussion, our approach to the system is the following combination of natural language processing and geospatial techniques: we collect and classify social media posts on Twitter related to vaccination; then categorize these posts by their sentiment, location and topic; then interpret the topics related to vaccine refusal and hesitancy; then spatially join and aggregate the results. This process enables evaluation of our survey hypotheses using the spatial topic clusters, as well as spatial examination of new discussions as it can be re-run over time. What follows are descriptions of each of the system's sub-processes.

### Methods – Data

The data related to vaccines for our context, also described briefly by Dredze et al.<sup>9</sup>, are Twitter posts (tweets) from the USA that we began collecting around the aforementioned measles outbreak in Disneyland. The system collects the data and classifies for relevance and sentiment. Initially we flagged tweets by keyword using the Twitter Streaming API<sup>i</sup>, specifying more than fifty hand-chosen keywords such as 'vaccine', 'shot', and 'immunization' similar to and validated by common practice<sup>2,6,31</sup>.

**Table 1:** Keywords used to filter Twitter data

vaccine,vaccines,shot,mmr,tdap,flushot,hpv,polio,rotavirus,chickenpox,smallpox,hepatitis,hepa,hepb,dtap,meningitis,shingles,vaccinate,vaccinated,vaccine,vaccines,vaccine,vacines,tetanus,diphtheria,pertussis,whoopingcough,dtp,dtwp,chickenpox,measles,mumps,rubella,varicella,diphtheria,haemophilus,papillomavirus,meningococcal,pneumococcal,rabies,tuberculosis,typhoid,yellowfever,immunizations,immunization,imunization,immune,imune,cholera,globulin,encephalitis,lyme
--

The Twitter Streaming API selects all tweets based on a given search (or a random 1% sample if a size threshold is exceeded); we obtained all matching tweets in the US during this time period, on the order of millions of tweets. See Dredze et al., for further details on these data.<sup>9</sup>

### Methods – Relevance and Sentiment

We trained supervised machine learning algorithms that, as part of the system, automatically classify tweets for relevance to vaccination and for sentiment, as sentiment analysis produces a measure of the expressed opinion in messages<sup>21</sup>. We obtained labeled training data using Amazon Mechanical Turk<sup>ii</sup> on randomly chosen subsets of our tweets to 1) tag them as being relevant to the topic of vaccines or not; 2) of those relevant, randomly choose and tag as having sentiment toward vaccines (neutral or non-neutral); 3) and of those that bear sentiment, randomly choose and tag as having positive or negative sentiment. While training the classifiers we conducted cross-validation and maximized

<sup>i</sup> <https://dev.twitter.com/streaming/overview>

<sup>ii</sup> <https://www.mturk.com/mturk/welcome>

precision and recall given tunable parameters.<sup>iii</sup> See Dredze et al., for further details on these classifiers.<sup>9</sup> Given these classifiers and the ability to run them over any tweets (our dataset and those incoming in real-time), we thus have the first part of the system, namely a real-time categorization of vaccine-related Twitter posts: those relevant to vaccines; of those relevant, those that bear sentiment; and of those that bear sentiment, their sentiment polarity.

### **Methods – Location Classification**

The next part of the system involves location classification. We use the Carmen Geolocation Toolkit<sup>10</sup> to automatically classify a tweet's location; such geolocation has been shown to be appropriate and effective in other public health studies<sup>10</sup>. Carmen improves upon information provided by the Streaming API, and returns location information at the country, state, county, and latitude/longitude levels if able.

### **Methods – Topic Classification**

The system uses topic modeling to determine the content discussed in our relevant tweets. Latent Dirichlet Allocation (LDA) is a commonly-used machine learning algorithm that automatically determines topics in collections of text<sup>1</sup>, common practice to automatically extract patterns and groups in text collections, of which social media data is a prime example. LDA assumes words in documents co-locate near other words (possibly across documents) because they are related, and the algorithm collects and reports groups of such related words, with the groups representing topics. Using the MACHine Learning for LanguagE Toolkit (MALLET)<sup>18</sup>, the system involves running LDA over our tweet dataset (the documents labeled as relevant to vaccination) to evaluate topics relevant to vaccine hesitancy. This produces an overall list of topics, and a parameterization of each tweet by topic (which is roughly proportional to relative composition by topic). Note that LDA is unsupervised; in general there is no guarantee that the algorithm will return a specific topic, and it is up to the analyst to determine topics' relevance and substance by analyzing the words and groups returned<sup>3</sup>. We leverage public health researchers' domain expertise to make such determinations. We note that in our context, LDA will show relevant topics because we have collected and categorized the tweets to fit a specific meta-topic (that of vaccines). By contrast, the substance of the relevant topics will be outputs of the system enabling hypothesis testing of our ground truth factors and exploration beyond.

### **Methods – Joins and Aggregations**

The system enables nonspatial aggregation and analysis on the tweets by sentiment and topic. More central to this paper, however: the tweets also have location data, which one may spatially join and aggregate using ArcMap (version 10.3), part of ArcGIS. ArcGIS is a geographic information system software that can generate maps of aggregated data and can calculate and display spatial statistics on those maps. One such statistic is the Getis-Ord  $G_i^*$  statistic for hotspot analysis<sup>14</sup>, valuable to the system because it indicates statistically significant high (low) point data if a point and its neighbors are high (low) in terms of some common variable.<sup>iv</sup> Using these maps and statistics, one may spatially analyze where vaccine tweets (our point data) occur, where sentiment occurs, and where topics occur (our common variables), with notions of how often they occur and whether statistically significant differences exist. Accordingly, the system provides a geographic result to accompany our topic-substance result concerning the survey results.

## **Results**

Running the topic model over all tweets in our dataset, we obtained information about topics and their distribution over our tweets. The system also produces classification results for each tweet in terms of its relevance, sentiment, and location. We may use a tweet's ID (e.g. "532385146419560448") to link its sentiment, location, and topic distribution. Given locations of relevant messages, we may filter by classification category and weight by topic distribution to find hotspots for discussion of a given discussion.

### **Results – Topics**

---

<sup>iii</sup> Relevance classifier (recall .91, precision .96); if relevant, whether contains sentiment about vaccines (recall .28, precision .63); if contains sentiment, is it positive vs negative (recall .85, precision .75). We chose to maximize precision in the second case because we were relying on the precision of our results in the positive/negative classifier. Such low recall is not an issue given the size of our dataset.

<sup>iv</sup> The definition of 'neighbor' is variable; what is appropriate depends highly on the input data. Some of many possible options for our topic proportions and tweet data are k-nearest-neighbors (weighting influence such that all points have k neighbors) or weighting influence based on inverse Euclidian. We chose the former due to ease of interpretation and calculation.

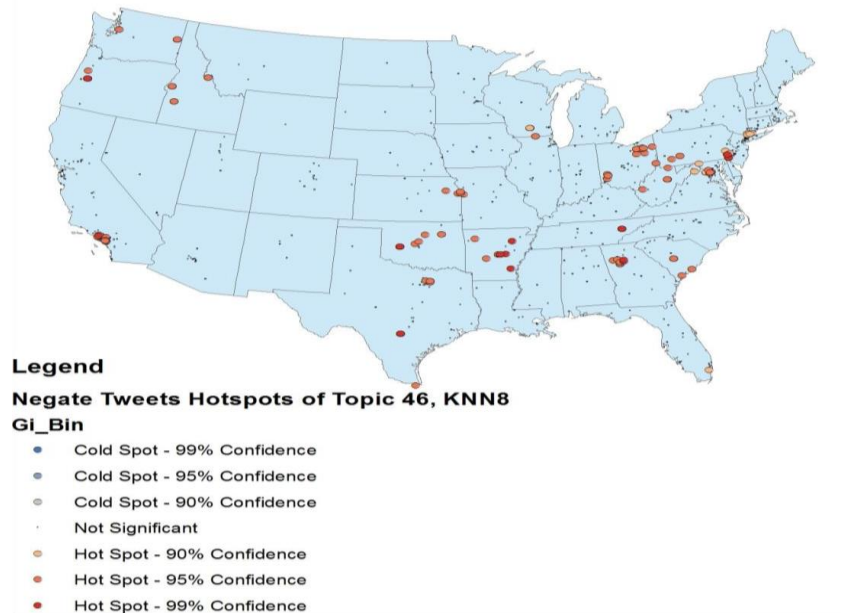
Specifically, the topic information consists of a relative weighting parameter for each topic for each tweet (roughly proportional to the proportion of each topic in the tweet), so one can get messages most representative of each topic. We ran the topic model on all messages, filtered by the regular expression *\*vacc\** to prune irrelevant / noisy topics in advance, and qualitatively interpreted the topics. Needing to specify the number of topics, we chose 50 to capture enough variability in our large dataset. As a proof of concept, we considered topic 46 in our further analyses. Topic 46 pertains to the California government's bill eliminating exemptions from vaccinations in schoolchildren. Below are example messages from this topic; our domain experts who performed identification and validation looked at both the tokens in the topic and representative messages when doing so, as is good practice<sup>5</sup>.

- “california governor signs strict school vaccine legislation gov jerry brown signs california bill imposing...”
- “jim carrey brands governor 'fascist' over vaccine law jim carrey called california gov jerry brown”
- “ahf criticizes dumb amp dumber star jim carrey for calling gov brown a fascist”
- “calif gov jerry brown launching frosted mercury flakes children's cereal to accompany vaccine mandate”

We chose this topic for two reasons: it is an arguably prominent anti-vaccination discussion in our data, and it is pertinent to a hypothesis validated by Quinn's previous work that “public trust / trust in government” affects such attitudes about medical decisions as vaccination, a common thread for validation. The analysis steps are the same regardless of topic chosen.

### Results – Hotspots for Topics

To identify hotbeds of these vaccine hesitancy discussions, we used the “Hot Spot Analysis” tool in ArcMap, which calculates the  $G_i^*$  statistic. We continued the proof of concept by considering only the contiguous United States, but the analysis is identical using different geographical boundaries (e.g. an individual state or a different country). We also limited our hot spot analysis only to the tweet messages classified as having negative sentiment about vaccines since our chosen topic was 46. As the definition of a neighborhood may vary depending on input data, we chose to spatially weight our input data via the k-nearest-neighbor (KNN) algorithm (using the default value of 8 neighbors suggested by ArcMap) to elegantly allow for such variations. This yielded the following map.



*Figure 1: Hotspots of the proportion of discussion of Topic 46 in the contiguous USA*

The hot-spot map of topic 46 shows statistically significant areas in the contiguous USA where the highest proportion of discussion of topic 46 is occurring in negative-sentiment vaccine messages on Twitter. For example, topic 46 is often discussed near LA and in the northern Appalachian region, among other areas. Such maps may be created for any permutation of classification and topic, and would yield any statistically significant results to be found among the

spatial data for each permutation. Note that this statistic does not merely highlight points that contain a lot of messages, but highlights points with statistically significant differences of message totals compared to neighboring points. Such significant results would (and do in the case of topic 46) yield convergent findings with survey data. Future work will more rigorously relate and apply this mixed methods approach.

## **Discussion**

The results outlined above yielded statistically significant geographic hot and cold spots in terms of individual topics in negative-sentiment vaccine messages on Twitter as a proof of concept. Such hotspots in a topic correspond to a discussion being statistically prevalent, and more prevalent in certain areas than others. That discussions pertinent to the trust in government results from Quinn's surveys (topic 46) are statistically significant in the first place both validates our approach and supports Quinn's findings on a larger scale. The fact that no significant cold spots are found among the topic 46 negative-sentiment map also validates our approach, as one would expect only hotspots in such topics pertaining to anti-vaccine discussions. This proof of concept showed that social media contains valuable information that is more granular and available more cheaply and quickly than through traditional survey methods. With further refinement, this information may be leveraged to replicate and compare with survey results.

In addition, such hot spot information is immediately actionable from a public health perspective, a valuable quality in the context of vaccine hesitancy. For example, one may target messages towards public policy think tanks in Arkansas to foster a more balanced approach to the debate about the government mandates on vaccination. Identification of such geolocated issues is valuable to public health officials as it provides low hanging fruit to address if interventions are known. For example, officials might value being able to reach all of Arkansas in a messaging campaign by only messaging Little Rock (if that were the only hotspot). The other side of the coin is also valuable, however, as evidence-based interventions may not yet exist. Officials may have been unaware of a specific geographic area and its opinions on a sub-issue of vaccine hesitancy, as hesitancy itself has been shown to vary across regions and within countries without a successful strategy.<sup>11</sup>

Thus the system's analysis of its real-time sentiment-topic data allowed us to identify individual discussions from the aggregate meta-topic, suggested the ability to verify survey hypotheses relating to those discussions, and suggested spatial targets for more effective use of public health resources. With expertise in both content and data analysis to fully understand and leverage the social media data, the system provides a promising opportunity to monitor real-time views.

## **Discussion – Limitations**

However, this system and its underlying approach may be improved. For example, the ability of Carmen<sup>10</sup> to augment location information could be increased such that it identifies information at a more granular level in more messages. This would affect the geospatial hotspot analysis, as one could improve results by grouping by levels of granularity with more and better location information. In addition, this proof of concept topic analysis returned 50 topics, but sensitivity analysis on this number as up- or downsizing could reduce noise. Thirdly, the open debate of social media analysis applies as well: whether social media discussions are a valid and accurate proxy for the rationales of the population at large. This applies both in terms of users' demographics (see below) and in terms of the potential for fake users, which recent research may be used to filter<sup>4</sup>. Fourthly, one should be cognizant of the (limited but nonzero) amount of technical supervision required: the system requires computational capacity and server administration, and it requires creating machine learning classifiers.<sup>9</sup>

## **Discussion – Future Work**

An additional limitation is that the topic models in LDA are subjective; there are alternative models and means of interpretation associated with them that could be employed. Paul and Dredze created an elegant framework for supervised topic models<sup>17,23</sup>, which could be adapted to our system, that would return topics seeded by specific a priori values (i.e., those in Quinn's survey results). Such seeded topics would remove subjectivity of topic interpretation, quantifiably associating topics with pre-determined results. Secondly, LDA is merely a long-running industry standard; an alternative is Linguistic Inquiry and Word Count (LIWC)<sup>33</sup>. In contrast to LDA which returns words that are co-located, LIWC counts psychologically relevant words into categories, producing output along dimensions such as “negative emotion words” or “tentative language”. These categories and their relative frequencies paint a picture of how the word user(s) consider their subject matter, in this case discussions about vaccines. Using LIWC would provide an alternative viewpoint that may be more easily interpreted using the framework of psychology.

Another aim of future work involves more explicit relations to traditional survey methods. One immediate improvement would be to aggregate tweets by user, which will enable user demographic classification<sup>7,30</sup> and other user-level statistics such as comparisons to known outbreaks of disease or to news coverage. With this information, and analysis related to retweets and news mentions, one might operationalize survey questions to individuals as different slices of our dataset, which for example would allow exploring and validating if demographics are related to one's rationales and opinions, especially those opinions relating to trust in government, as previous work has suggested<sup>25,26</sup>. Aggregating information by user would also allow the system to further the question of whether social media may be used as a proxy for the population at large, both in terms of demographics and in terms of coverage of topic discussion. The representativeness of social media users is an open question, whether relating to pro- or anti-vaccination communities or to the population as a whole. The analyses in this paper combined with demographic classification would allow us to determine how representative our social media users are of our target population(s).

## Conclusions

Given the problem of tracking and understanding discussion in a population and the context of vaccine hesitancy, we have as a first step created a pipeline of natural language processing and geospatial techniques that enable real-time statistical analysis of different discussions in a population across space. This system showed statistically significant spatial hotspots of discussion in the USA that provide actionable insights for the time-sensitive context. Given the financial and computational ease of gathering and processing swaths of social media data, this system can be used to monitor real-time views, and, easily extensible, suggests the ability to verify traditional survey methods in broader spatial contexts.

## Acknowledgements

Thank you to Amelia Jamison for her helpful feedback and topic analysis.

Dr. Dredze has received consulting fees from Directing Medicine LLC and Sickweather LLC, who use social media for public health surveillance.

## References

1. Blei DM, Ng AY, Jordan MI. Latent Dirichlet Allocation. *J Mach Learn Res.* 2003 Mar;3:993–1022.
2. Broniatowski DA, Paul MJ, Dredze M. National and Local Influenza Surveillance through Twitter: An Analysis of the 2012–2013 Influenza Epidemic. *PLOS ONE.* 2013 Dec 9;8(12):e83672.
3. Chang J, Boyd-Graber JL, Gerrish S, Wang C, Blei DM. Reading tea leaves: How humans interpret topic models. In: *Nips [Internet].* 2009 [cited 2017 Mar 8]. p. 1–9. Available from: <https://papers.nips.cc/paper/3700-reading-tea-leaves-how-humans-interpret-topic-models.pdf>
4. Cheng J, Danescu-Niculescu-Mizil C, Leskovec J. Antisocial Behavior in Online Discussion Communities. *arXiv:150400680 [cs, stat] [Internet].* 2015 Apr 2 [cited 2016 Dec 8]; Available from: <http://arxiv.org/abs/1504.00680>
5. Chuang J, Manning CD, Heer J. Termite: Visualization Techniques for Assessing Textual Topic Models. In: *Proceedings of the International Working Conference on Advanced Visual Interfaces [Internet].* New York, NY, USA: ACM; 2012 [cited 2017 Mar 8]. p. 74–77. (AVI '12). Available from: <http://doi.acm.org/10.1145/2254556.2254572>
6. Conover MD, Goncalves B, Ratkiewicz J, Flammini A, Menczer F. Predicting the Political Alignment of Twitter Users. In: *2011 IEEE Third International Conference on Privacy, Security, Risk and Trust (PASSAT) and 2011 IEEE Third International Conference on Social Computing (SocialCom).* 2011. p. 192–9.
7. Culotta A, Ravi NK, Cutler J. Predicting Twitter User Demographics using Distant Supervision from Website Traffic Data. *Journal of Artificial Intelligence Research.* 2016;55:389–408.
8. Culotta A. Towards Detecting Influenza Epidemics by Analyzing Twitter Messages. In: *Proceedings of the First Workshop on Social Media Analytics [Internet].* New York, NY, USA: ACM; 2010 [cited 2016 Mar 10]. p. 115–122. (SOMA '10). Available from: <http://doi.acm.org/10.1145/1964858.1964874>
9. Dredze M, Broniatowski DA, Smith MC, Hilyard KM. Understanding Vaccine Refusal: Why We Need Social Media Now. *American Journal of Preventive Medicine.* 2016 Apr;50(4):550–2.
10. Dredze M, Paul MJ, Bergsma S, Tran H. Carmen: A twitter geolocation system with applications to public health. In: *AAAI Workshop on Expanding the Boundaries of Health Informatics Using AI (HIAI).* Citeseer; 2013. p. 20–24.
11. Dubé E, Gagnon D, MacDonald NE. Strategies intended to address vaccine hesitancy: Review of published reviews. *Vaccine.* 2015 Aug 14;33(34):4191–203.

12. Dubé E, Gagnon D, Nickels E, Jeram S, Schuster M. Mapping vaccine hesitancy—Country-specific characteristics of a global phenomenon. *Vaccine*. 2014 Nov 20;32(49):6649–54.
13. Freimuth VS, Musa D, Hilyard K, Quinn SC, Kim K. Trust during the early stages of the 2009 H1N1 pandemic. *Journal of health communication*. 2014;19(3):321–339.
14. Getis A, Ord JK. The Analysis of Spatial Association by Use of Distance Statistics. *Geographical Analysis*. 1992 Jul 1;24(3):189–206.
15. Halsey NA, Salmon DA. Measles at Disneyland, a Problem for All Ages Measles at Disneyland. *Ann Intern Med*. 2015 May 5;162(9):655–6.
16. MacDonald NE. Vaccine hesitancy: Definition, scope and determinants. *Vaccine*. 2015 Aug 14;33(34):4161–4.
17. McAuliffe JD, Blei DM. Supervised Topic Models. In: Platt JC, Koller D, Singer Y, Roweis ST, editors. *Advances in Neural Information Processing Systems 20* [Internet]. Curran Associates, Inc.; 2008 [cited 2016 Mar 17]. p. 121–128. Available from: <http://papers.nips.cc/paper/3328-supervised-topic-models.pdf>
18. McCallum AK. MALLETT: A Machine Learning for Language Toolkit [Internet]. 2002. Available from: <http://mallet.cs.umass.edu>
19. Mollema L, Harmsen IA, Broekhuizen E, Clijnk R, De Melker H, Paulussen T, et al. Disease Detection or Public Opinion Reflection? Content Analysis of Tweets, Other Social Media, and Online Newspapers During the Measles Outbreak in the Netherlands in 2013. *J Med Internet Res* [Internet]. 2015 May 26 [cited 2016 Mar 4];17(5). Available from: <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC4468573/>
20. Omer SB, Salmon DA, Orenstein WA, deHart MP, Halsey N. Vaccine Refusal, Mandatory Immunization, and the Risks of Vaccine-Preventable Diseases. *New England Journal of Medicine*. 2009 May 7;360(19):1981–8.
21. Pang B, Lee L. Opinion Mining and Sentiment Analysis. *Found Trends Inf Retr*. 2008 Jan;2(1–2):1–135.
22. Paul MJ, Dredze M, Broniatowski D. Twitter Improves Influenza Forecasting. *PLoS Currents* [Internet]. 2014 [cited 2016 Mar 5]; Available from: <http://currents.plos.org/outbreaks/?p=39911>
23. Paul MJ, Dredze M. SPRITE: Generalizing Topic Models with Structured Priors. *Transactions of the Association for Computational Linguistics*. 2015 Jan 20;3(0):43–57.
24. Quinn SC, Hilyard K, Castaneda-Angarita N, Freimuth VS. Public acceptance of peramivir during the 2009 H1N1 influenza pandemic: implications for other drugs or vaccines under emergency use authorizations. *Disaster Med Public Health Prep*. 2015 Apr;9(2):166–74.
25. Quinn SC, Kumar S, Freimuth VS, Kidwell K, Musa D. Public willingness to take a vaccine or drug under Emergency Use Authorization during the 2009 H1N1 pandemic. *Biosecurity and bioterrorism: biodefense strategy, practice, and science*. 2009;7(3):275–290.
26. Quinn SC, Parmer J, Freimuth VS, Hilyard KM, Musa D, Kim KH. Exploring communication, trust in government, and vaccination intention later in the 2009 H1N1 pandemic: results of a national survey. *Biosecurity and bioterrorism: biodefense strategy, practice, and science*. 2013;11(2):96–106.
27. Quinn SC, Thomas T, Kumar S. The Anthrax Vaccine and Research: Reactions from Postal Workers and Public Health Professionals. *Biosecurity and bioterrorism*. 2008 Dec;6(4):321–33.
28. Roush SW, Murphy TV, Vaccine-Preventable Disease Table Working Group a. Historical comparisons of morbidity and mortality for vaccine-preventable diseases in the United States. *JAMA*. 2007 Nov 14;298(18):2155–63.
29. Salathé M, Freifeld CC, Mearns SR, Tomasulo AF, Brownstein JS. Influenza A (H7N9) and the Importance of Digital Epidemiology. *New England Journal of Medicine*. 2013 Aug 1;369(5):401–4.
30. Sap M, Park G, Eichstaedt JC, Kern ML, Stillwell DJ, Kosinski M, et al. Developing Age and Gender Predictive Lexica over Social Media. In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)* [Internet]. Association for Computational Linguistics; 2014 [cited 2016 Jun 8]. p. 1146–1151. Available from: <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.672.9851>
31. Signorini A, Segre AM, Polgreen PM. The Use of Twitter to Track Levels of Disease Activity and Public Concern in the U.S. during the Influenza A H1N1 Pandemic. *PLOS ONE*. 2011 May 4;6(5):e19467.
32. Smith MC, Broniatowski DA, Paul MJ, Dredze M. Towards Real-Time Measurement of Public Epidemic Awareness: Monitoring Influenza Awareness through Twitter. In: *AAAI Spring Symposium on Observational Studies through Social Media and Other Human-Generated Content*. Stanford, CA; 2016.
33. Tausczik YR, Pennebaker JW. The Psychological Meaning of Words: LIWC and Computerized Text Analysis Methods. *Journal of Language and Social Psychology*. 2010 Mar 1;29(1):24–54.
34. Winter K, Carol Glaser, James Watt, Kathleen Harriman. Pertussis Epidemic — California, 2014 [Internet]. 2014 [cited 2017 Mar 8]. Available from: <https://www.cdc.gov/mmwr/preview/mmwrhtml/mm6348a2.htm>