

SOCIAL MEDIA MINING FOR PUBLIC HEALTH MONITORING AND SURVEILLANCE

MICHAEL J. PAUL,^{1*} ABEED SARKER,² JOHN S. BROWNSTEIN,³ AZADEH NIKFARJAM,²
MATTHEW SCOTCH,² KAREN L. SMITH,⁴ GRACIELA GONZALEZ²

¹ *Department of Information Science, University of Colorado, Boulder, CO 80304, USA*

² *Department of Biomedical Informatics, Arizona State University, Tempe, AZ 85287, USA*

³ *Department of Pediatrics, Harvard Medical School, Boston MA 02115, USA*

⁴ *School of Pharmacy, Regis University, Denver, CO 80221, USA*

**E-mail: michael.j.paul@colorado.edu*

This paper describes topics pertaining to the session, “Social Media Mining for Public Health Monitoring and Surveillance,” at the Pacific Symposium on Biocomputing (PSB) 2016. In addition to summarizing the content of the session, this paper also surveys recent research on using social media data to study public health. The survey is organized into sections describing recent progress in public health problems, computational methods, and social implications.

Keywords: Social media; data mining; natural language processing; public health.

1. Background

Social media platforms have seen unprecedented worldwide growth. For example, as of June 30, 2015, Twitter has over 300 million active monthly users, 77% of whom are outside of the US.¹ Social networks form a platform for people to share and discuss their views and opinions, and many share their health-related information both in general-purpose social media (such as Twitter, Facebook or Instagram) and in health-related social networks (communities focusing specifically on health issues, such as DailyStrength or MedHelp). Advances in automated data processing, machine learning and natural language processing (NLP) present the possibility of utilizing these massive data sources for public health monitoring and surveillance, as long as researchers are able to address the methodological challenges unique to this media.

Numerous studies have been published recently in this realm, including studies on pharmacovigilance,² identifying smoking cessation patterns,³ identifying user social circles with common experiences (like drug abuse),⁴ monitoring malpractice,⁵ and tracking infectious disease spread.^{6–8} A systematic review⁹ conducted in 2014 found numerous attempts to use this user-generated data, but none yet integrated in national surveillance programs, noting the promise and challenges of the field quite succinctly:

“More direct access to such [social media] data could enable surveillance epidemiologists to detect potential public health threats such as rare, new diseases or early-level warnings for epidemics. But how useful are data from social media and the Internet, and what is the potential to enhance surveillance? The challenges of using these emerging surveillance systems for infectious disease epidemiology, including the specific resources needed, technical requirements, and acceptability to public health practitioners and policymakers, have wide-reaching implications for public health surveillance in the 21st century.”⁹

The use of social media for health monitoring and surveillance indeed has many drawbacks

and difficulties, particularly if done automatically. For example, traditional NLP methods that are applied to longer texts have proven to be inadequate when applied to short texts, such as those found in Twitter.² Something seemingly simple, such as searching and collecting relevant postings, has also proven to be quite challenging, given the amount of data and the diverse styles and wording used by people to refer to the topic of interest in colloquial terms (semantic heterogeneity) inherent to this type of media.

The goal of this session was to attract researchers that have explored automatic methods for the collection, extraction, representation, analysis, and validation of social media data for public health surveillance and monitoring, including epidemiological and behavioral studies. It serves as a unique forum to discuss novel approaches to text and data mining methods that respond to the specific requirements of social media and that can prove invaluable for public health surveillance. Research topics presented at this session include:

- Early detection of disease outbreaks¹⁰
- Medication safety, including drug interactions¹¹ and dietary supplement safety¹²
- Health behaviors, including diet success¹³ and smoking cessation¹⁴
- Individual well-being,¹⁵ which affects mental and physical health

This paper first summarizes the current state of, and recent advances in, social media mining for health monitoring, focusing on examples of promising research areas (Section 2), technical challenges (Section 3), and societal implications and considerations (Section 4). We then provide an overview of the research presented at this session in Section 5, with concluding remarks in Section 6.

2. Expanding the Frontiers of Public Health

We begin by summarizing recent research in some key areas of public health for which social media mining has been especially popular and fruitful, with an emphasis on how these focus areas are evolving to increase public health impact.

2.1. *Disease Surveillance: Beyond Influenza*

Disease surveillance is one of the longest-running use cases for social media mining. Some of the earliest work using web data for public health surveillance was to estimate influenza prevalence from search query volumes.¹⁶ This idea was made famous with Google's widely-used Flu Trends service.^{17,18} Google Flu Trends recently ended their service (as of August 2015), but Google will continue to share their data with academic research labs.¹⁹ While search queries were the original data sources for web-based disease surveillance, social media has since become a popular data source for influenza monitoring, including weblogs²⁰ and microblogs, especially Twitter.^{21–25}

Influenza has been by far the most commonly surveilled disease, in part due to its widespread prevalence—it affects millions of people each year (causing 3,000–50,000 yearly deaths in the US²⁶), making it both an important disease to monitor and a disease that is widely discussed in social media. The original motivation for using web data to estimate influenza prevalence is that it can be estimated in real-time, in contrast to traditional gov-

ernment systems—the national surveillance coordinated by the Centers for Disease Control and Prevention in the US, for example, is one to two weeks out of date. However, it has been argued that social media-based influenza surveillance has limited utility in many scenarios, as many agencies and institutions already conduct timely influenza surveillance.²⁷

More recently, web-based disease surveillance research has moved in new directions with potentially higher impact:

Other infectious diseases More recent social media research has considered disease surveillance for infectious diseases other than influenza. For example, a number of researchers have used search and tweet data to track dengue fever.^{28–31} Others have used Twitter to monitor cholera,⁷ *E. coli*,³² and ebola.^{33,34}

Forecasting While most early work on web-based disease surveillance focused on estimating the current week disease prevalence (referred to as “nowcasting”), more recent work has attempted to *forecast* disease prevalence, using web data to predict prevalence weeks into the future.^{35–38} The ability to accurately predict future levels of disease prevalence will greatly help with planning and preparedness.

High-impact locations Much work with influenza has focused on surveillance at the national level in countries such as the US, but more recent work has focused on locations that would benefit more from real-time surveillance: countries with fewer existing surveillance resources³⁹ and fine-grained locations, such as hospitals^{40,41} and mass gatherings.⁴²

2.2. *Pharmacovigilance*

Pharmacovigilance, which primarily involves the monitoring of adverse reactions caused by medications, is another established use case of social media.⁴³ Users discuss their health-related experiences, including the use of prescription drugs, side effects and treatments on social media, which makes social networks unique and robust sources of information about health, drugs and treatments. Research has focused on the detection of user posts mentioning adverse reactions and the extraction of drug-adverse reaction association signals, utilizing data from specialized health communities and forums,^{2,44–46} online reviews of drugs⁴⁷ and generic networks such as Twitter.^{48–51}

Adverse reaction detection A number of studies focus on the automatic classification of user posts to determine if adverse reactions are mentioned. Common approaches involve utilizing annotated data sets to perform supervised classification to identify adverse reaction assertive posts and/or personal experiences of adverse reactions.^{46,48,52,53} Supervised classification approaches require manually annotated data and recent advances in research have seen the creation of such data sets.^{52–54} One important challenge that has been frequently discussed in supervised learning tasks is the data imbalance in social media data.^{52,53}

Discovering drug-adverse reaction associations Some research has concentrated on extracting specific adverse reaction mentions (and their lexical variants) and identifying as-

sociations between specific drugs and adverse reactions. Most past approaches are lexicon-based^{2,45,55} and recent approaches have applied supervised learning techniques for extraction.⁵¹ Following the extraction of concepts, co-occurrence metrics have been applied for quantifying drug-adverse reaction associations.⁴⁴

2.3. Behavioral Medicine

Another rapidly expanding area of social media surveillance is understanding behaviors that affect health, such as smoking and diet. It has been argued that behavioral medicine will play a prominent role in the digital surveillance revolution, because there is a large knowledge gap in many areas of behavioral medicine.⁵⁶ We summarize recent research in a few key areas.

Smoking and substance abuse One of the major uses of social media to study behavioral medicine has been to understand smoking and tobacco use.⁵⁷ Social media can be used to understand availability of and interest in various nicotine and tobacco products,⁵⁸ including electronic cigarettes, which are a rapidly evolving market for which social media has provided much faster intelligence than traditional sources.^{59–61} Social networks have also been analyzed to understand smoking cessation and online social support for cessation.^{3,62–64}

Other substance abuse issues have been studied as well, including trends in alcohol use^{65,66} and problem drinking.^{67,68} Some researchers have focused on using social media data for monitoring prescription drug abuse.^{4,69–73} Specialized social networks have been used for analyzing the effects of drug reformulation⁷⁴ or the phases of drug abuse recovery.⁷⁵ Among generic social networks, Twitter is becoming increasingly popular for monitoring patterns of specific prescription medication abuse.^{4,69}

Diet and fitness A number of researchers have analyzed food consumption patterns in Instagram^{76,77} and Twitter,^{78,79} including seasonal patterns in weight loss.⁸⁰ Researchers have also studied physical activities in Twitter,^{81,82} including measuring outcomes of fitness goals.⁸³

3. Technical Challenges of Social Media Mining

There are a number of challenges with automated text analysis, particularly when working with data from social media. We describe some of the key analytic tasks needed for public health mining, along with recent advances in these technologies.

3.1. Processing Informal Text

A key challenge with automatic data mining of social media is that standard NLP tools, which are traditionally trained on formal text (e.g., newswire), do not adapt well to the informal, non-standard language used online. Some researchers have created NLP tools, such as part-of-speech taggers and named entity recognizers, specifically for Twitter.^{84,85} This can help researchers apply NLP to tweets, although this is not a general-purpose solution: tools tailored to Twitter may not work well on other social media platforms.

A particular challenge in the domain of health is that laypeople on social media may not use accurate medical terminology. One solution to this issue is to analyze text that mentions

common symptoms, rather than references to specific illnesses.⁸⁶ There has also been research to correct and normalize medical terminology,⁸⁷ and there is a large body of research on general language normalization for social media text.⁸⁸

3.2. *Sentiment Analysis*

A particular branch of NLP that shows promise for social media monitoring is *sentiment analysis*.⁸⁹ Sentiment analysis involves automatically ascribing positive, negative, or neutral sentiment to portions of text that express opinions. Sentiment analysis has been applied to social media in interesting ways to understand important public health issues. We provide a few examples here.

Sentiment analysis has been used to understand public attitudes toward vaccination by analyzing Twitter messages.⁹⁰ For example, the study in Ref. 91 found that negative sentiment toward vaccines spreads through social networks more than positive sentiment. Sentiment has also been analyzed in the context of drug abuse, in order to understand public interest in drugs. For example, researchers have measured shifts in public attitudes toward marijuana.^{92,93} Sentiment analysis is particularly applicable to online reviews, which have been analyzed for public health in the domain of online doctor and healthcare provider reviews, to understand patient perceptions of care quality.^{94,95} The studies in Refs. 96,97 found that sentiment inferred from reviews is significantly correlated with existing provider quality metrics.

However, sentiment analysis does not work as well for short text, such as tweets.⁹⁸ Sentiment classification is an active area of NLP research, and improvements in this technology will lead to improvements in understanding public opinion and awareness.

3.3. *Richer Language Understanding*

Much of the research on social media mining for health monitoring has used relatively simple methods of text analysis, such as dictionary associations. While simple approaches can work reasonably well, there is an upper limit to their performance, and future improvements will require NLP tools that can extract richer meaning from text.

Richer NLP can even improve seemingly simple tasks. For example, while early research showed that tweets with keywords such as “flu” are well-correlated with influenza prevalence,²¹ more recent research has shown that flu is discussed in different ways on Twitter, for example, whether a user is describing a personal experience or simply sharing news of the flu season, and whether a user is personally sick or whether they are describing a family member or co-worker.^{25,99} These distinctions can affect the performance of influenza surveillance, and such distinctions require NLP systems that incorporate richer *n*-gram and linguistic features.²⁵

Richer NLP techniques have also been applied to concept extraction tasks, such as adverse drug reaction mention extraction. Early techniques primarily focused on lexicon-based approaches, where the natural language mentions of the elements of interest are encoded in lexicons and these are utilized to detect their mentions in text.^{2,45,46,55} These techniques have led to the development of health-related lexical resources from social media sources (e.g., the Consumer Health Vocabulary¹⁰⁰). The use of colloquial language, however, limits the performance of such approaches. With the creation of annotated data in recent years, supervised

machine learning approaches are becoming increasingly popular, and they have also shown promising performance in quantitative evaluations.^{51,101}

4. Societal Implications and Considerations

There are a number of social and societal implications of using social media for public health. We briefly discuss some key considerations here.

4.1. *Impact of Social Media Monitoring*

There is currently a gap between what is possible with social media monitoring—which many studies have demonstrated successfully, as described in Section 2—with what is being done in practice. As noted in the review in Ref. 9, existing social media systems have not widely integrated with national surveillance. However, the landscape is beginning to change. The US government has expressed interest in using social media for health surveillance, with both the CDC and the Department of Health and Human Services (HHS) soliciting submissions of systems that monitor social media for health issues.^{102,103} Private companies, such as Sickweather, make social media monitoring available to the general public.

One hurdle in bringing social media monitoring to practice is gaining trust of practitioners and the public. For example, trust in web-based disease surveillance was eroded after researchers showed significant failings of the popular Google Flu Trends system.¹⁰⁴ More time will be needed to understand how such systems perform in practice. In the meantime, researchers must validate their social media models carefully to ensure progress is being made.¹⁰⁵

4.2. *Ethics of Social Media Research*

There are a number of ethical considerations to keep in mind when using social media data for health research. One of the key concerns hinges on the extent to which social media data should be treated as public versus private data.¹⁰⁶ Even though social media data are publicly available, social media users may not intend or wish for their data to be used for research.¹⁰⁷ Users may not be aware that their social media data is publicly available,¹⁰⁸ and may have expectations of privacy even in public settings.¹⁰⁹ The distinction between public and private data becomes additionally complicated by the fact that machine learning algorithms can make inferences about private attributes, even if not explicitly stated in public data.¹¹⁰

Addressing these issues involves an ongoing conversation among Internet researchers,¹¹¹ and a number of scholars have written about using big data for research.¹¹² For more discussion of social media ethics in public health research, see Refs. 113–115.

5. Session Overview

This session hosted cutting-edge research in many of the public health areas described in Section 2. We briefly summarize the contributions below.

5.1. *Disease Surveillance*

Ofoghi *et al.*¹⁰ presented research on disease-related emotion detection in tweets, suggesting that emotion tweets can be utilized to detect and monitor disease outbreaks. This work intro-

duced NLP classifiers to categorize tweets into various *emotions* (e.g., “anger”, “surprise”). The distributions of emotions were then analyzed in datasets of tweets pertaining to the ebola epidemic in 2014–2015. The authors found that the distributions differed among tweets at the time and place of an outbreak compared to outside tweets. These results suggest that emotion classification could help distinguish outbreak-related tweets from other disease discussion.

This research is an example of using richer NLP models to categorize disease-related tweets in useful ways, as discussed in Section 3.3: it is not enough to know that a tweet discusses ebola, but rather *how* ebola is being discussed.

5.2. *Pharmacovigilance*

The session hosted two papers on pharmacovigilance.

Correia *et al.*¹¹ investigated the utility of Instagram—an increasingly active social media platform—as a source of information for adverse drug reactions (ADRs). Instagram constitutes a potentially novel data source, in contrast to most social media-based ADR research which has focused on platforms such as Twitter and Facebook. This study analyzed, and introduced visualization tools for, Instagram messages mentioning various drugs used for depression. The results show that health issues are commonly discussed on Instagram, and there is potential for identifying ADRs, including interactions with other drugs and products.

Sullivan *et al.*¹² focused on adverse reactions to dietary supplements, which are products that are not currently well-monitored. This study analyzed Amazon.com reviews of nutritional supplements, and used a topic modeling system to categorize products based on their potential danger, as suggested in reviews. In the study, the proposed automated system agreed with human annotators 69.4% of the time, suggesting that automated methods can potentially be used to flag dangerous products.

5.3. *Behavioral Medicine*

The session included multiple studies that fall broadly in the category of behavioral medicine.

Aphinyanaphongs *et al.*¹⁴ analyzed tweets for mentions of e-cigarette use. Because e-cigarettes constitute a relatively new product and public health phenomenon, real-time surveillance is needed to better understand usage patterns in the population. This work developed classifiers to identify tweets which mention e-cigarettes, as well as tweets which mention using e-cigarettes to support smoking cessation. The study developed a baseline classification performance of up to .90 AUC for detecting use and .94 AUC for detecting smoking cessation intent. The results show potential for measuring e-cigarette use from Twitter.

Weber and Achananuparp¹³ analyzed public food diaries from the application, MyFitnessPal, and constructed models to predict whether users will or will not meet their daily caloric goals. By analyzing the predictive features, this study provides insights into what features are predictive of diet success or failure. Some results are expected, such as oil and butter contributing to diet failure and fruits contributing to diet success, while some insights are non-trivial, such as differences between types of meat. Future work points to insights from more complex features, such as the interactions of dietary groups.

Schwartz *et al.*¹⁵ developed models to predict the state of well-being of individuals from

their Facebook data, where *well-being* reflects positive mood as well as additional constructs such as meaning in life and engagement in activities. Using *n*-gram and topic features, the authors built classifiers to estimate various metrics of well-being at the level of individual Facebook messages, as well as the aggregate level of a user's entire stream. The goal of such research is to improve our understanding of the determinants and consequences of well-being, which is correlated with outcomes of both mental and physical health.

6. Concluding Remarks

The goal of this session was to create a single venue for cross-disciplinary researchers to present research on social media mining for public health monitoring and surveillance. The session provided a forum to share new research in a variety of important public health areas, including the detection of disease outbreaks and awareness; pharmacovigilance, including interactions with natural products and dietary supplements; and various issues related to behavioral medicine, including weight loss, e-cigarette use, and well-being. Through these projects, researchers also advanced the technology needed to understand social media text, for example by developing new NLP classifiers, new topic model variations, and new visualization systems. Given the ever-increasing amount of social media data around the world, interest in such systems will only increase over time.

References

1. Twitter: Company Facts <https://about.twitter.com/company>.
2. R. Leaman and L. Wojtulewicz, Towards internet-age pharmacovigilance: extracting adverse drug reactions from user posts to health-related social networks, in *Proceedings of the 2010 Workshop on Biomedical Natural Language Processing*, 2010.
3. L. L. Struik and N. B. Baskerville, *J. Med. Internet Res.* **16**, p. e170 (2014).
4. L. C. Hanson, B. Cannon, S. Burton and C. Giraud-Carrier, *J Med Internet Res* **15**, p. e189 (September 2013).
5. A. Nakhasi, R. J. Passarella, S. G. Bell, M. J. Paul, M. Dredze and P. J. Pronovost, Malpractice and Malcontent: Analyzing Medical Complaints in Twitter, in *AAAI Fall Symposium on Information Retrieval and Knowledge Discovery in Biomedical Text*, 2012.
6. M. J. Paul and M. Dredze, You are what you Tweet: Analyzing Twitter for public health, in *International Conference on Weblogs and Social Media*, 2011.
7. R. Chunara, J. R. Andrews and J. S. Brownstein, *American Journal of Tropical Medicine and Hygiene* **86**, 39 (2012).
8. D. A. Broniatowski, M. J. Paul and M. Dredze, *PLoS ONE* **8**, p. e83672 (2013).
9. E. Velasco, T. Agheneza, K. Denecke, G. Kirchner and T. Eckmanns, *Milbank Q* **92**, 7 (Mar 2014).
10. B. Ofoghi, M. Mann and K. Verspoor, Towards early discovery of salient health threats: A social media emotion classification technique, in *PSB*, 2016.
11. R. B. Correia, L. Li and L. M. Rocha, Monitoring potential drug interactions via network analysis of Instagram user timelines, in *PSB*, 2016.
12. R. Sullivan, A. Sarkar, K. O'Connor, A. Goodin, M. Karlsrud and G. Gonzalez, Monitoring dietary supplements: Challenges and promises of mining user comments for adverse events, in *PSB*, 2016.

13. I. Weber and P. Achananuparp, Insights from machine-learned diet success prediction, in *PSB*, 2016.
14. Y. Aphinyanaphongs, A. Lulejian, D. P. Brown, R. Bonneau and P. Krebs, Classification for automatic detection of e-cigarette use and use for smoking cessation from twitter: a feasibility pilot, in *PSB*, 2016.
15. H. A. Schwartz, M. Sap, M. L. Kern, J. C. Eichstaedt, A. Kapelner, M. Agrawal, E. Blanco, L. Dziurzynski, G. Park and L. H. Ungar, Predicting individual well-being through the language of social media, in *PSB*, 2016.
16. G. Eysenbach, Infodemiology: tracking flu-related searches on the web for syndromic surveillance (2006).
17. J. Ginsberg, M. H. Mohebbi, R. S. Patel, L. Brammer, M. S. Smolinski and L. Brilliant, *Nature* **457**, 1012 (2009).
18. S. Cook, C. Conrad, A. L. Fowlkes and M. H. Mohebbi, *PLoS ONE* **6**, p. e23610 (2011).
19. Google Flu Trends, The next chapter for Flu Trends <http://googleresearch.blogspot.com/2015/08/the-next-chapter-for-flu-trends.html>, (2015), Accessed 2015-08-29.
20. C. D. Corley, D. J. Cook, A. R. Mikler and K. P. Singh, *International Journal of Environmental Research and Public Health* **7**, 596 (2010).
21. A. Culotta, Towards detecting influenza epidemics by analyzing Twitter messages, in *ACM Workshop on Soc.Med. Analytics*, 2010.
22. A. Signorini, A. M. Segre and P. M. Polgreen, *PLoS ONE* **6**, p. e19467 (2011).
23. E. Aramaki, S. Maskawa and M. Morita, Twitter catches the flu : Detecting influenza epidemics using Twitter the university of tokyo, in *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, 2011.
24. V. Lampos and N. Cristianini, *ACM Trans. Intell. Syst. Technol.* **3**, 1 (2012).
25. A. Lamb, M. J. Paul and M. Dredze, *Proceedings of NAACL-HLT 2013*, 789 (2013).
26. Estimating Seasonal Influenza-Associated Deaths in the United States: CDC Study Confirms Variability of Flu http://www.cdc.gov/flu/about/disease/us_flu-related_deaths.htm.
27. D. R. Olson, K. J. Konty, M. Paladini, C. Viboud and L. Simonsen, *PLoS Computational Biology* **9** (2013).
28. E. H. Chan, V. Sahai, C. Conrad and J. S. Brownstein, *PLoS Neglected Tropical Diseases* **5** (2011).
29. B. M. Althouse, Y. Y. Ng and D. A. T. Cummings, *PLoS Neglected Tropical Diseases* **5** (2011).
30. R. T. Gluskin, M. A. Johansson, M. Santillana and J. S. Brownstein, *PLoS Negl Trop Dis* **8**, p. e2713 (February 2014).
31. J. Gomide, A. Veloso, W. Meira, Jr., V. Almeida, F. Benevenuto, F. Ferraz and M. Teixeira, Dengue surveillance based on a computational model of spatio-temporal locality of Twitter, in *Proceedings of the 3rd International Web Science Conference*, 2011.
32. E. Diaz-Aviles and A. Stewart, Tracking Twitter for epidemic intelligence: Case study: Ehec/hus outbreak in germany, 2011, in *Proceedings of the 4th Annual ACM Web Science Conference*, 2012.
33. M. Odlum and S. Yoon, *Am J Infect Control* **43**, 563 (Jun 2015).
34. M. Odlum, How Twitter can support early warning systems in ebola outbreak surveillance, in *Annual Meeting of the American Public Health Association*, 2015.
35. A. F. Dugas, M. Jalalpour, Y. Gel, S. Levin, F. Torcaso, T. Igusa and R. E. Rothman, *PLoS ONE* **8**, p. e56176 (2013).
36. J. Shaman, A. Karspeck, W. Yang, J. Tamerius and M. Lipsitch, *Nat Commun* **4**, p. 2837 (2013).
37. E. Nsoesie, M. Marathe and J. Brownstein, *PLoS Currents* **5**, 1 (2013).
38. M. J. Paul, M. Dredze and D. Broniatowski, *PLOS Currents Outbreaks* (2014).

39. M. Paul, M. Dredze, D. Broniatowski and N. Generous, Worldwide influenza surveillance through Twitter, in *AAAI Workshop on the World Wide Web and Public Health Intelligence*, 2015.
40. D. A. Broniatowski, M. Dredze, J. M. Paul and A. Dugas, *JMIR Public Health Surveill* **1**, p. e5 (2015).
41. O. M. Araz, D. Bentley and R. L. Muelleman, *Am J Emerg Med* **32**, 1016 (Sep 2014).
42. E. Yom-Tov, D. Borsa, I. J. Cox and R. A. McKendry, *J. Med. Internet Res.* **16**, p. e154 (2014).
43. A. Sarker, R. Ginn, A. Nikfarjam, K. O'Connor, K. Smith, S. Jayaraman, T. Upadhaya and G. Gonzalez, *Journal of Biomedical Informatics* **54**, 202 (2015).
44. A. Nikfarjam and G. H. Gonzalez, Pattern mining for extraction of mentions of adverse drug reactions from user comments, in *AMIA Annual Symposium*, 2011.
45. A. Benton, L. Ungar, S. Hill, S. Hennessy, J. Mao, A. Chung, C. E. Leonard and J. H. Holmes, *J Biomed Inform* **44**, 989 (Dec 2011).
46. A. Yates, N. Coharian and O. Frieder, Extracting adverse drug reactions from forum posts and linking them to drugs, in *SIGIR Workshop on Health and Discovery*, 2013.
47. A. Yates and N. Goharian, ADRTrace: Detecting expected and unexpected adverse drug reactions from user reviews on social media sites, in *Proceedings of the 35th European Conference on Advances in Information Retrieval*, (Berlin, Heidelberg, 2013).
48. J. Bian, U. Topaloglu and F. Yu, Towards large-scale Twitter mining for drug-related adverse events, in *International Workshop on Smart Health and Wellbeing*, 2012.
49. K. O'Connor, P. Pimpalkhute, A. Nikfarjam, R. Ginn, K. L. Smith and G. Gonzalez, *AMIA Annu Symp Proc* **2014**, 924 (2014).
50. C. C. Freifeld, J. S. Brownstein, C. M. Menone, W. Bao, R. Filice, T. Kass-Hout and N. Dasgupta, *Drug Saf* **37**, 343 (May 2014).
51. A. Nikfarjam, A. Sarker, K. O'Connor, R. Ginn and G. Gonzalez, *J Am Med Inform Assoc* (March 2015).
52. R. Ginn, P. Pimpalkhute, A. Nikfarjam, A. Patki, K. O'Connor, A. Sarker and G. Gonzalez, Mining Twitter for adverse drug reaction mentions: a corpus and classification benchmark, in *Proceedings of the Fourth Workshop on Building and Evaluating Resources for Health and Biomedical Text Processing (BIOTXTM)*, 2014.
53. A. Sarker and G. Gonzalez, *Journal of Biomedical Informatics* **53**, 196 (2014).
54. I. Segura-Bedmar, R. revert and P. Martinez, Detecting drugs and adverse events from spanish health social media streams, in *Proceedings of the 5th international workshop on health text mining and information analysis (LOUHI)*, 2014.
55. S. Yeleswarapu, A. Rao, T. Joseph, V. G. Sapradeep and R. Srinivasan, *BMC Medical Informatics and Decision Making* **14** (2014).
56. J. W. Ayers, B. M. Althouse and M. Dredze, *JAMA* **311**, 1399 (2014).
57. K. W. Prier, M. S. Smith, C. Giraud-Carrier and C. L. Hanson, Identifying health-related topics on Twitter: An exploration of tobacco-related tweets as a test topic, in *Proceedings of the 4th International Conference on Social Computing, Behavioral-cultural Modeling and Prediction*, SBP'11 (Springer-Verlag, 2011).
58. M. Myslin, S. H. Zhu, W. Chapman and M. Conway, *J. Med. Internet Res.* **15**, p. e174 (2013).
59. J. W. Ayers, K. M. Ribisl and J. S. Brownstein, *American Journal of Preventive Medicine* **40**, 448 (2011).
60. J. Huang, R. Kornfield, G. Szczypka and S. L. Emery, *Tob Control* **23 Suppl 3**, 26 (July 2014).
61. H. Cole-Lewis, A. Varghese, A. Sanders, M. Schwarz, J. Pugatch and E. Augustson, *J. Med. Internet Res.* **17**, p. e208 (2015).
62. N. K. Cobb, A. L. Graham, M. J. Byron, D. B. Abrams and Workshop Participants, *Journal of Medical Internet Research* **13** (2011).

63. M. Rocheleau, R. S. Sadasivam, K. Baquis, H. Stahl, R. L. Kinney, S. L. Pagoto and T. K. Houston, *J. Med. Internet Res.* **17**, p. e18 (2015).
64. J. J. Prochaska, C. Pechmann, R. Kim and J. M. Leonhardt, *Tob Control* **21**, 447 (Jul 2012).
65. J. H. West, P. C. Hall, K. Prier, C. L. Hanson, C. Giraud-Carrier, E. S. Neeley and M. D. Barnes, *Open Journal of Preventative Medicine* **2** (2012).
66. Y. Aphinyanaphongs, B. Ray, A. Statnikov and P. Krebs, Text classification for automatic detection of alcohol use-related tweets, in *International Workshop on Issues and Challenges in Social Computing*, 2014.
67. A. K. Fournier and S. W. Clarke, *Journal of Psychosocial Research on Cyberspace* **5** (2011).
68. M. A. Moreno, D. A. Christakis, K. G. Egan, L. N. Brockman and T. Becker, Associations between displayed alcohol references on facebook and problem drinking among college students (2012).
69. C. L. Hanson, S. H. Burton, C. Giraud-Carrier, J. H. West, M. D. Barnes and B. Hansen, *J. Med. Internet Res.* **15**, p. e62 (2013).
70. N. Genes and M. Chary, Twitter discussions of nonmedical prescription drug use correlate with federal survey data, in *Medicine 2.0 Conference*, 2014.
71. B. Chan, A. Lopez and U. Sarkar, *PLoS ONE* **10**, p. e0135072 (08 2015).
72. T. K. Mackey, B. A. Liang and S. A. Strathdee, *J. Med. Internet Res.* **15**, p. e143 (2013).
73. P. M. Coloma, B. Becker, M. C. J. M. Sturkenboom, E. M. van Mulligen and J. A. Kors, *Drug Safety* **38**, 921 (2015).
74. E. C. McNaughton, P. M. Coplan, R. A. Black, S. E. Weber, H. D. Chilcoat and B. S. F., *Journal of Medical Internet Research* **16** (May 2014).
75. D. MacLean, S. Gupta, A. Lembke, C. Manning and J. Heer, Forum77: An analysis of an online health forum dedicated to addiction recovery, in *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, CSCW '15 (ACM, New York, NY, USA, 2015).
76. S. Sharma and M. De Choudhury, Detecting and Characterizing Nutritional Information of Food and Ingestion Content in Instagram, in *WWW*, 2015.
77. Y. Mejova, H. Haddadi, A. Noulas and I. Weber, #foodporn: Obesity patterns in culinary interactions, in *Proceedings of the 5th International Conference on Digital Health 2015*, 2015.
78. S. Abbar, Y. Mejova and I. Weber, *CoRR* **abs/1412.4361** (2014).
79. D. Fried, M. Surdeanu, S. Kobourov, M. Hingle and D. Bell, Analyzing the language of food on social media, in *IEEE International Conference on Big Data*, 2014.
80. G. M. Turner-McGrievy and M. W. Beets, *Transl Behav Med* **5**, 160 (Jun 2015).
81. N. Zhang, S. Campo, K. F. Janz, P. Eckler, J. Yang, L. G. Snetselaar and A. Signorini, *Journal of medical Internet research* **15** (2013).
82. V. L. D. Reis and A. Culotta, Using matched samples to estimate the effects of exercise on mental health from Twitter, in *AAAI*, 2015.
83. E. Kiciman and M. Richardson, Towards decision support and goal achievement: Identifying action-outcome relationships from social media, in *KDD*, 2015.
84. K. Gimpel, N. Schneider, B. O'Connor, D. Das, D. Mills, J. Eisenstein, M. Heilman, D. Yogatama, J. Flanigan and N. A. Smith, Part-of-speech tagging for Twitter: Annotation, features, and experiments, in *Association for Computational Linguistics (ACL)*, 2011.
85. A. Ritter, S. Clark, Mausam and O. Etzioni, Named entity recognition in tweets: An experimental study, in *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2011.
86. F. Gesualdo, G. Stilo, E. Agricola, M. V. Gonfiantini, E. Pandolfi, P. Velardi and A. E. Tozzi, *PLoS ONE* **8**, p. e82489 (2013).
87. L. Nie, M. Akbari, T. Li and T.-S. Chua, A joint local-global approach for medical terminology

- assignment, in *SIGIR 2014 Workshop on Medical Information Retrieval*, 2014.
88. J. Eisenstein, What to do about bad language on the internet, in *Proceedings of the North American Chapter of the Association for Computational Linguistics (NAACL)*, 2013.
 89. B. Pang and L. Lee, *Foundations and Trends in Information Retrieval* **2**, 1 (2008).
 90. M. Salathe and S. Khandelwal, *PLoS Comput Biol* **7**, p. e1002199 (2011).
 91. M. Salathé, D. Q. Vu, S. Khandelwal and D. R. Hunter, *EPJ Data Science* **2** (2013).
 92. P. A. Cavazos-Rehg, M. Krauss, S. L. Fisher, P. Salyer, R. A. Gruzca and L. J. Bierut, *J Adolesc Health* **56**, 139 (Feb 2015).
 93. L. Thompson, F. P. Rivara and J. M. Whitehill, *Cyberpsychology, Behavior, and Social Networking* **18**, 311 (2015).
 94. A. López, A. Detz, N. Ratanawongsa and U. Sarkar, *J Gen Intern Med* **27**, 685 (Jun 2012).
 95. S. Brody and N. Elhadad, Detecting salient aspects in online reviews of health providers, in *AMIA Annual Symposium*, 2010.
 96. J. Segal, M. Sacopulos, V. Sheets, I. Thurston, K. Brooks and R. Puccia, *J. Med. Internet Res.* **14**, p. e50 (2012).
 97. B. C. Wallace, M. J. Paul, U. Sarkar, T. A. Trikalinos and M. Dredze, *Journal of the American Medical Informatics Association (JAMIA)* (2014).
 98. A. Agarwal, B. Xie, I. Vovsha, O. Rambow and R. Passonneau, Sentiment analysis of Twitter data, in *Proceedings of the Workshop on Languages in Social Media*, 2011.
 99. R. Nagar, Q. Yuan, C. C. Freifeld, M. Santillana, A. Nojima, R. Chunara and J. S. Brownstein, *J. Med. Internet Res.* **16**, p. e236 (2014).
 100. Consumer Health Vocabulary <http://consumerhealthvocab.org/>.
 101. H. Sampathkumar, X. wen Chen and B. Luo, *BMC Medical Informatics and Decision Making* **14** (October 2014).
 102. Now trending <https://nowtrending.hhs.gov/>, Accessed: 2015-08-23.
 103. C. for Disease Control and Prevention, Predict the influenza season challenge <https://www.federalregister.gov/articles/2013/11/25/2013-28198/announcement-of-requirements-and-registration-for-the-predict-the-influenza-season-challenge>, (2013).
 104. D. Lazer, R. Kennedy, G. King and A. Vespignani, *Science* **343**, 1203 (2014).
 105. T. Bodnar and M. Salathé, Validating models for disease detection using Twitter, in *Proceedings of the 22Nd International Conference on World Wide Web Companion*, 2013.
 106. R. McKee, *Health policy (Amsterdam, Netherlands)* **110**, 298 (2013).
 107. J. M. Hudson and A. Bruckman, *Inf. Soc.* **20**, 127 (2004).
 108. Y. Liu, K. P. Gummadi, B. Krishnamurthy and A. Mislove, Analyzing facebook privacy settings: User expectations vs. reality, in *Proceedings of the 2011 ACM SIGCOMM Conference on Internet Measurement Conference*, 2011.
 109. J. C. H. Bromseth, Public places – public activities? methodological approaches and ethical dilemmas in research on computer-mediated communication, in *Researching ICTs in Context*, ed. A. Morrison (Inter/Media Report, 2002) pp. 33–61.
 110. E. Horvitz and D. Mulligan, *Science* **349**, 253 (Jul 2015).
 111. C. Fiesler, A. Young, T. Peyton, A. S. Bruckman, M. Gray, J. Hancock and W. Lutters, Ethics for studying online sociotechnical systems in a big data world, in *Proceedings of the 18th ACM Conference Companion on Computer Supported Cooperative Work & Social Computing*, 2015.
 112. D. boyd and K. Crawford, *Information, Communication & Society* **15** (2012).
 113. C. M. Rivers and B. L. Lewis, *F1000Research* **3** (2014).
 114. M. Conway, *Journal of medical Internet research* **16** (2014).
 115. E. Vayena, M. Salathé, L. C. Madoff and J. S. Brownstein, *PLoS Comput Biol* **11**, p. e1003904 (2015).