Published online: 16.08.2019

Recent Advances in Using Natural Language Processing to Address Public Health Research Questions Using Social Media and Consumer-Generated Data

Mike Conway, Mengke Hu, Wendy W. Chapman
Department of Biomedical Informatics, University of Utah, Salt Lake City, Utah, United States

Summary

Objective: We present a narrative review of recent work on the utilisation of Natural Language Processing (NLP) for the analysis of social media (including online health communities) specifically for public health applications.

Methods: We conducted a literature review of NLP research that utilised social media or online consumer-generated text for public health applications, focussing on the years 2016 to 2018. Papers were identified in several ways, including PubMed searches and the inspection of recent conference proceedings from the Association of Computational Linguistics (ACL), the Conference on Human Factors in Computing Systems (CHI), and the International AAAI (Association for the Advancement of Artificial Intelligence) Conference on Web and Social Media (ICWSM). Popular data sources included Twitter, Reddit, various online health communities, and Facebook.

Results: In the recent past, communicable diseases (e.g., influenza, dengue) have been the focus of much social media-based NLP health research. However, mental health and substance use and abuse (including the use of tobacco, alcohol, marijuana, and opioids) have been the subject of an increasing volume of research in the 2016 - 2018 period. Associated with this trend, the use of lexicon-based methods remains popular given the availability of psychologically validated lexical resources suitable for mental health and substance abuse research. Finally, we found that in the period under review "modern" machine learning methods (i.e. deep neural-network-based methods), while increasing in popularity, remain less widely used than "classical" machine learning methods.

Keywords

Natural Language Processing; text mining; social media; public health

Yearb Med Inform 2019:208-17
http://dx.doi.org/10.1055/s-0039-1677918

1 Introduction

Social media is a valuable source of data for public health research. It is estimated that 75% of Internet users have read or watched online health information content, and 26% of Internet users have posted (or shared) their personal health information online [1]. This large-scale sharing of health information makes social media and Online Health Communities (OHC) a valuable and abundant source of data for addressing public health questions. Social media — including online consumer generated OHC data — provide a ready source of timely, abundant data that can serve as a valuable resource for several broad types of public health applications, including surveillance, health communication, sentiment analysis, and understanding the natural history of a disease, injury, or health behaviour. Research on utilising social media in conjunction with Natural Language Processing (NLP) for public health applications is a robust and growing area of study, with dedicated meetings[1] and a now well-established research community [2]. Regarding surveillance, the importance of mental health and substance abuse surveillance is increasingly recognised [3]. This growth is unsurprising given that it is estimated that mental health and substance abuse constitute approximately 10.4% of the global burden of disease and are the leading cause of years lived with disability, imposing direct and indirect costs on the world economy of around US$2.5 trillion [4]. The study of health communication is another area of research that uses social media in conjunction with NLP methods, particularly in the area of understanding and quantifying vaccine hesitancy and refusal. NLP can support public health researchers in identifying common health-related misconceptions, and in turn, devising more effective health communication methods [5]. Similarly, sentiment analysis with respect to products relevant to public health (e.g. marijuana-related products, e-cigarettes) and the health behaviours that they facilitate is a further area of research [6]. Finally, social media provide a valuable data source for studies focussed on understanding and analysing the natural history of a disease, illness or injury, especially in the context of new and re-emerging diseases and rapid changes in health behaviour [7].

The key changes we have observed since 2016 — apart from the growth in research related to mental health and substance abuse and the increasing interest in "modern" machine learning methods — include a move towards integrating social media analysis with the Electronic Health Record (EHR) [8], in part as a means of obtaining valuable diagnostic "ground truth". A further shift of note is the increased interest in elucidating ethical issues in the application of NLP (and machine learning more generally) to social media for public health applications, particularly with respect to protecting the rights of those users suffering from potentially stigmatising conditions [9].

[1] For example, the Social Media Mining for Health Applications (SMM4H) Workshop or the Computational Linguistics and Clinical Psychology (CLPsych) Workshop

Challenges in developing high performance NLP methods for social media have been extensively enumerated, but in summary, major outstanding problems include the use of non-standard grammar, the use of rapidly changing and often non-standard slang terms , spelling variation in informal consumer-generated text, the rapidly changing nature of social media language, and finally the identification (and filtering) of jokes, memes, and advertising [2].

In this paper, we review literature from the period 2016-2018 regarding the application of NLP methods to social media data as a means of addressing public health research questions, focussing specifically on new application areas and the adoption of new methods. A distinctive feature of this review is an emphasis on the increasing volume of research focussed on ethics-related issues involved in using consumer-generated data for public health research.

2 Methods

Our paper selection process involved the following steps. First, we searched PubMed, the Association for Computational Linguistics Anthology, the Proceedings of the Conference on Human Factors in Computer Systems (CHI), and the Proceedings of the International AAAI (Association for the Advancement of Artificial Intelligence) Conference on Web and Social Media (ICWSM) using a variety of social media and NLP-related keywords. Second, we manually inspected Tables of Contents for the *Journal of the American Medical Informatics Association*, the *Journal of Biomedical Informatics*, and the *Journal of Medical Internet Research*. In this first pass, over 1,800 papers were identified. After reviewing abstracts, we reduced the number of papers reviewed to 130. In order to increase the tractability of the reviewing task, we further winnowed the papers to 71. This winnowing process was designed to capture a large swathe of both application areas and methods, and cannot be interpreted as a comment on the quality of research.

Only the papers that both demonstrated a clear public health focus and explicitly utilised NLP or text mining methods were retained. Papers that reported on the results of qualitative content analysis or professional standards for health communication using social media without reference to NLP were excluded. Papers that discussed ethical issues pertaining to the use of social media for public health applications and research were retained. References dated outside the period 2016-2018 have been included in order to provide important context. The use of these references does not imply that they form part of the document set defined by the inclusion criteria.

The papers reviewed utilise social media from several different sources, including Twitter, Reddit, Weibo, Facebook, and online discussion forums (see Figure 1 and Tables 1 & 2).

The vast majority of the papers reviewed focussed on analysing English language text (68 papers), with two papers focussing on Chinese text [76, 77] and one paper focussing on Japanese text [31]. With respect to the geographical location of first authors, most of the articles emerged from North America (55), with Europe (7), and Asia (including Australasia and Turkey) (6) all represented.

The reviewed papers can be grouped into several health-related categories, including vaccine hesitancy and refusal, communicable diseases surveillance (including sexually transmitted infections, [STIs]), cancer, substance abuse, pharmacovigilance, and mental health (see Table 2). A wide range of methods were used, including "classical" machine learning (e.g., Random Forests, Support Vector Machines [SVM]), "modern" machine learning (e.g., Convolutional Neural Networks [CNN], Recurrent Neural Networks [RNN][2]), and lexicon-based approaches). Among the lexicon-based approaches, the Linguistic Inquiry and Word Count (LIWC) lexicon, a dictionary of words arranged into numerous psychological dimensions, is used extensively in many of the papers reviewed, especially in the areas of mental health and substance abuse [79].

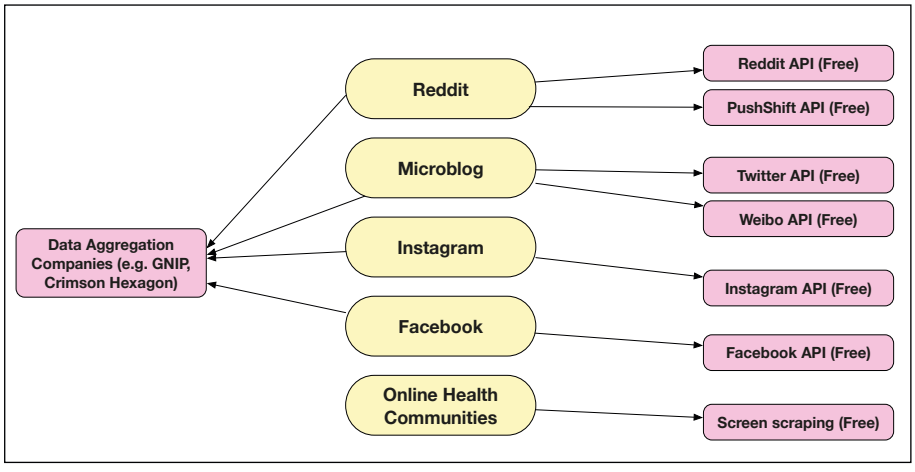

Fig. 1 Social media data sources. Note that this list is not exhaustive.

3 Results

3.1 Vaccine Hesitancy and Refusal

Vaccine hesitancy — defined by the World Health Organisation as referring to a "delay in acceptance or refusal of vaccines despite availability of vaccination services"[3] — has been a growing subject of research during

[2] Note that the terms "classical" and "modern" machine learning are, from a historical perspective, misnomers, given the roots of neural network theory in the mid-twentieth century [78].

[3] https://www.who.int/immunization/programmes_systems/vaccine_hesitancy/en/

Table 1 Number of papers by topic and data source. Note that papers can occur in several categories

Data Source	Vac[a]	Comm[b]	Cancer[c]	SA[d]	Pharmaco[e]	STI[f]	MH[g]	Total
Reddit	-	1	-	3	-	1	13	18
Twitter	3	3	1	17	7	1	9	41
Instagram	-	-	-	-	-	-	1	1
Facebook	1	-	-	-	-	-	3	4
OHC[h]	1	-	2	2	1	-	6	12
Weibo	-	1	-	-	-	-	1	2
WhatsApp	-	-	-	1	-	-	-	1
Youtube	-	-	-	1	-	-	-	1
Yik-Yak	-	-	-	1	-	-	-	1
Tumblr	-	-	-	-	-	-	1	1

[a]Vaccination hesitancy and refusal; [b]Health communication; [c]Cancer; [d]Substance Abuse; [e]Pharmacovigilance; [f]Sexually transmitted infections; [g]Mental health; [h]Online Health Communities

Table 2 Data Sources and Topics [Note that ethics-related papers are excluded from this table as they are frequently concerned with social media in general.]

Data Source	Vac[a]	Comm[b]	Cancer	SA[c]	Pharmaco[d]	STI[e]	MH[f]
Reddit	-	[10]	-	[11-13]	-	[14]	[15-27]
Twitter	[28-30]	[31-33]	[34]	[6, 12, 35-49]	[50-56]	[57]	[18, 58-65]
Instagram	-	-	-	-	-	-	[18]
Facebook	[66]	-	-	-	-	-	[8, 18, 67]
OHC[g]	[5]	-	[68, 69]	[12, 13]	[50]	-	[70-75]
Weibo	-	[32]	-	-	-	-	[76]
Tumblr	-	-	-	-	-	-	[18]

[a]Vaccination hesitancy and refusal; [b]Communicable diseases; [c]Substance Abuse; [d]Pharmacovigilance; [e]Sexually transmitted infections; [f]Mental health; [g]Online Health Communities

the review period, with NLP methods applied to social media data in an attempt to develop insights into how best to understand and improve health communication as well as quantifying the degree of vaccine hesitancy in a community.

Of the five papers reviewed in this section (see Table 3), three utilised Twitter data [29-30], one utilised Facebook data [66], and one further paper utilised data derived from an online health community, in this case mothering.com [5]. Supervised machine learning [30] and unsupervised machine learning [5, 28, 29] were both represented. Three of the papers reviewed used classical machine learning methods [5, 29, 30], and one used modern machine learning methods [30], with surveillance [28-30], health communication [5, 28-30, 66], and sentiment analysis [28-30, 66], all frequently studied topics. The LIWC lexicon has been used either to characterise public attitudes towards vaccination in general [66], or as a tool to explore the purported link between autism and the Measles, Mumps, and Rubella vaccine [28]. This last study aimed at investigating key differences between users who are longstanding vaccination advocates, long standing anti-vaccination advocates, or users who had recently adopted an anti-vaccination orientation. Vaccination to protect against the Human Papillomavirus Virus (HPV) — a vaccine typically administered to adolescent boys and girls to prevent future sexual transmission of the disease — was also the subject of reviewed research, with high performance sentiment classifiers developed (AUC: 0.92) [30], and LDA (Latent Dirichlet Allocation) topic modeling used to identify a number of vaccine-hesitancy-related topics, including clinical evidence and vaccination harms [29].

In a further example of novel research, Tangherlini et al., produced a statistical-mechanical network model representing relationships between "actants" (actors) that is used to automatically extract typical narratives and "story fragments" related to vaccination issues, evidencing a narrative framework related to a pronounced distrust of government and medical authority [5].

3.2 Communicable Diseases and Sexually Transmitted Infections

Systems designed to use social media data for pandemic public health surveillance have existed for almost 13 years [80, 81], and approaches that are variously referred to as infodemiology [82], digital disease detection [83], and digital epidemiology [84] are by now well established, particularly for dengue, influenza, and more recently, ebola. In addition, significant research efforts have centered on the study of STI, despite some methodological concerns regarding the willingness of users with STIs to disclose their status on social media.

In order to investigate the changing prevalence of a number of health related topics, Park et al., [10] observed that ebola discussions were characterised by concerns about risks and symptoms, while influenza was associated with terms like "CDC" and "H1N1". Another study focussed on influenza misdiagnoses [33], achieving an F-score of 0.76. Regarding STIs, one study demonstrated statistically significant associations between Twitter data from 2012 and official Centers for Disease Control syphilis prevalence data from 2013 [57], with a related study discovering that the most frequent STIs discussed were intermediate (non-reportable) STIs like genital herpes

and HPV, with more serious (reportable) diseases like syphilis and gonorrhoea discussed less frequently [14].

Of the six papers reviewed (see Table 4), four used Twitter data [31-33, 57], and two used Reddit data [10, 14], while Al-Garadi et al., provided a review that concentrated on Twitter and Weibo, the Chinese language microblog service [32]. Two of the papers reviewed described the use of supervised machine learning methods [31, 32], three papers used unsupervised machine learning methods [10, 14, 32], and one used a lexicon-based approach [57]. Machine learning methods were used to perform a variety of tasks, including surveillance [10, 14, 31-33, 57], health communication [32], and sentiment analysis [32]. Several studies concentrated on influenza surveillance using English [10, 33] and Japanese [31] Twitter data.

Table 3 Summary of vaccine-related papers

Data Source	SML[a]	UML[b]	UML[b]	CML[c]	MML[d]	Surv[e]	HC[f]	Senti[g]	Lexicon[h]
Twitter	[30]	[28, 29]	[28, 29]	[29, 30]	[30]	[28-30]	[28-30]	[28-30]	[28]
Facebook	-	-	-	-	-	-	[66]	[66]	[66]
OHC[i]	-	[5]	[5]	[5]	-	-	[5]	-	-

[a]Supervised machine learning (e.g., Support Vector Machines, Random Forests); [b]Unsupervised machine learning (e.g., Latent Dirichlet Allocation, K-means); [c]Classical machine learning (e.g., Random Forests, Support Vector Machines); [d]Modern machine learning (e.g., Convolutional Neural Networks); [e]Surveillance; [f]Health communication; [g]Sentiment analysis; [h]Lexicon-based methods; [i]Online health communities

Table 4 Summary of communicable diseases and STI-related papers

Data Source	SML[a]	UML[b]	CML[c]	MML[d]	Surv[e]	HC[f]	Senti[g]	Lexicon[h]
Reddit	-	[10, 14]	[10, 14]	-	[10, 14]	-	-	-
Twitter	[31, 32]	[32]	[31-33]	-	[31-33, 57]	[32]	[32]	[57]
Weibo	[32]	[32]	[32]	-	[32]	[32]	[32]	-

[a]Supervised machine learning; [b]Unsupervised machine learning; [c]Classical machine learning; [d]Modern machine learning; [e]Surveillance; [f]Health communication; [g]Sentiment analysis; [h]Lexicon-based methods

3.3 Cancer

Work on using NLP and text-mining methods to understand issues directly related to cancer (diagnosis, treatment, and management) are less well developed than some of the other areas considered in this review (e.g., mental health and substance abuse). Of the three cancer-related papers reviewed (see Table 5), one utilised Twitter data [34], and two utilised data derived from an online health community [68, 69]. All the papers discussed used both classical and modern machine learning methods, with modern machine learning methods performing better than classical machine learning methods, albeit by a narrow margin in the case of Zhang et al.'s work on identifying chemotherapy-related Twitter accounts by account type [34]. Zhang et al., observed that Twitter accounts belonging to individuals focussed on "personal chemotherapy experience and emotions", whereas professional accounts typically provided a neutral presentation of chemotherapy side effects [34]. Two of the papers were centred on health communication, broadly conceived [68, 69], with one paper focusing on sentiment analysis [34]. Concentrating specifically on the patient experience of breast cancer, one study [68] aimed at characterizing how forum topics changed over time depending on the individual's time since diagnosis and cancer state, and found that diagnosis is the most frequent class in the early stages of cancer treatment, with diagnosis (and treatment) related discussions declining over the course of a user's cancer journey.

3.4 Substance Abuse

This section is concerned with reviewing work centred on the use of social media, in conjunction with NLP methods, to address substance abuse research questions, focussing on *opioid abuse*, *tobacco, e-cigarette and marijuana use*, and *alcohol abuse*. Interesting work on drug abuse — particularly new and emerging products — is increasingly evident in the literature. NLP methods are needed to deal with ambiguity and colloquial expressions used on social media (such as "bath salts", "kitty cat", or "miaow miaow" for mephedrone [44]).

Of the twenty-two papers discussed in this section, three are focussed on opioid abuse [35, 41, 42], eight on tobacco and marijuana use [6, 12, 13, 40, 43, 45, 46, 49], one on alcohol abuse [36], and one on the street drug, mephedrone [44]. Twitter is the most popular source of data (18 papers) [6, 11, 12, 35-49], with Reddit [11-13], and online health communities [12, 13], both represented. Supervised machine learning (8 papers – all utilising Twitter data) and unsupervised machine learning (11 papers) were both evident in the reviewed papers, with classical machine learning approaches more common than modern neural-network-based approaches (17 and 2 papers, respectively). Two of the papers reviewed utilized a rule-based approach. Table 6 summarises the reviewed substance abuse-related papers.

3.4.1 Opioid Abuse

Opioid abuse is now recognised as one of the leading public health problems in the United States[4], and an important — albeit slightly less pressing — concern in many developed and developing countries. The crisis in the US is due to historical changes in drug prescription policies and practices that have encouraged both the licit and illicit use of highly addictive opioid-based painkillers[5].

[4] https://www.cdc.gov/drugoverdose/epidemic/index.html

[5] https://www.drugabuse.gov/drugs-abuse/opioids/opioid-overdose-crisis

Every year in the United States, over 72,000 people die as a direct consequence of using opioids[6], making the need to understand emerging opioid-related behaviours and user trajectories especially pressing. One study concentrated on identifying public reactions to the opioid epidemic by identifying the most popular opioid-related topics tweeted by users [41]. Topics identified included discussions related to the possibility of promoting marijuana as a substitute for opioids, discussions related to the growing opioid market in North America, and discussions related to news reports advocating the use of buprenorphine — a narcotic used to treat opioid addiction — for adolescents experiencing opioid use disorders. Another study [35] aimed at detecting marketing and sale of opioids by illicit online sellers. The authors observed that the frequency of tweets directly related to illegal activity was relatively low when compared with other kinds of opioid mentions. A similar observation was made for tweets promoting the illegal online sale of fentanyl [42]. In this context, unsupervised approaches are of significant value for understanding changes in a rapidly developing online environment.

3.4.2 Tobacco, E-Cigarette, and Marijuana Use and Abuse

Tobacco use is declining in popularity in much of the developed world (the proportion of smokers in the US has declined by over half since 1964 and now stands at 16.8% among adults, and approximately half that among high school students [85]). However, despite this decrease in tobacco use, there has been a dramatic increase — now plateauing — in the use of e-cigarettes since their introduction to developed world markets in around 2007 [86]. This increase has occurred in the context of a lack of consensus regarding both the safety of the product [87] and its potential efficacy as a smoking cessation device [88]. In addition to these shifts in tobacco use, there have also been substantial changes in the regulation of marijuana products, particularly in the US context, and these changes have led — it has been suggested [89] — to an increase in marijuana use [90]. Given these public health concerns, using NLP to investigate tobacco, e-cigarette, and marijuana use, has become an active research area, especially to classify discussions [6, 12, 43, 45, 46] or to determine whether a particular user is above or below 21 years of age [40]. Reported findings included evidence that Twitter users frequently discussed ways in which e-cigarettes can be used in the workplace in a bid to circumvent smoking bans [43], and evidence that hookah was discussed more frequently at the weekend, indicating its use is associated with leisure activities, while reported tobacco use tends to be more consistent across the week [40]. In addition, authors observed that different social media services manifested distinctly different cultures regarding e-cigarette use, e.g., sensory experiences vs. psychological factors associated with quitting [13]. Rule-based approaches were used to identify *where* people reported using e-cigarettes, with 39% of posts referring to e-cigarette use in the classroom [49]. Other studies aimed at describing strategies for marketing Little Cigars & Cigarillos (LCC) and observed that 83% of identified LCC tweets referred to marijuana, and 29% of LCC tweets referenced memes [45].

[6] https://www.cdc.gov/drugoverdose/data/statedeaths.html

Table 5 Summary of cancer-related papers

Data Source	SML[a]	UML[b]	CML[c]	MML[d]	Surv[e]	HC[f]	Senti[g]	Lexicon[h]
Twitter	[34]	[34]	[34]	[34]	-	-	[34]	-
OHC[i]	[68, 69]	[68]	[68, 69]	[68, 69]	-	[68, 69]	-	-

[a]Supervised machine learning; [b]Unsupervised machine learning; [c]Classical machine learning; [d]Modern machine learning; [e]Surveillance; [f]Health communication; [g]Sentiment analysis; [h]Lexicon; [i]Online Health Communities

Table 6 Summary of substance abuse-related papers

Data source	SML[a]	UML[b]	CML[c]	MML[d]	Surv[e]	HC[f]	Senti[g]	Lexicon[h]
Reddit	-	[11-13]	[11-13]	-	[12]	-	-	[13]
Twitter	[6, 36, 40, 45-49]	[6, 12, 35,37, 39, 41, 42, 43, 45]	[6, 12, 35, 36, 38-43, 45-49]	[6, 37]	[11, 12, 35, 36, 38, 39, 42, 44, 47-49]	[43]	[46-48]	[44]
OHC[i]	-	[12, 13]	[12, 13]	-	[12]	-	-	[13]

[a]Supervised machine learning; [b]Unsupervised machine learning; [c]Classical machine learning; [d]Modern machine learning; [e]Surveillance; [f]Health communication; [g]Sentiment analysis; [h]Lexicon; [i]Online Health Communities

3.4.3 Alcohol Abuse

Alcohol abuse was the seventh leading risk-factor worldwide for both death and disability in 2016. In the same year, among males aged 15-49, alcohol was a causal factor in 12% of deaths [91]. One of the reviewed studies [36] yielded the surprising result that— in the US at least — a positive correlation exists between excessive county-level alcohol consumption and higher education, suggesting that highly educated counties drink more, or at least tweet more about their drinking.

3.5 Pharmacovigilance

Pharmacovigilance — i.e. the post-market surveillance of drugs — was an early health-related focus for social media NLP [92, 93] and has remained an important subject of research, with applications including the identification of mentions of Adverse

Drug Reactions (ADRs) [51, 55]. One recent study focussed on topics related to Thyroid Hormone Replacement Therapy (THRT), particularly on the identification of side effects [50]. It was discovered that male and female users of THRT had different experiences and concerns regarding side effects, with women primarily concerned about the effect of the drug on personal appearance and men more concerned about potential pain symptoms associated with the drug.

A recent significant development in pharmacovigilance research was the instigation of the SMM4 2017 shared task. The shared task consisted of three subtasks: automatic identification of ADRs, automatic classification of tweets that explicitly mentioned medication consumption, and normalization of ADR mentions. Important outputs of this effort included a publicly available corpus [51] and language models [55] for future research. In addition to this work on ADR identification and normalization, the identification of semantic relationships — chiefly *causal* relationships — between drug and symptom mentions had been a focus of research [52, 53]. A key challenge associated with this task is the difficulty involved in distinguishing between drug use as a response to a particular symptom ("I have a horrible headache and just took some ibuprofen") and the existence of a symptom as a side effect of a drug ("Ever since I started taking Sertraline I've felt like crap"). Despite the difficulty of this task, Bollegala et al., achieved a moderately high F-score (0.74) using a skip-gram based method [52].

Six of the pharmacovigilance papers reviewed used Twitter as a data source [51-56], while one used an online health community (see Table 7). Four of the papers used supervised methods [51-54] and five used unsupervised methods [50, 53-56] with five using classical machine learning methods [50-53, 56] and three using modern machine learning methods [51, 54, 55], with (unsurprisingly given the topic of pharmacovigilance) surveillance being the main application area.

Table 7 Summary of pharmacovigilance-related papers

Data Source	SML[a]	UML[b]	CML[c]	MML[d]	Surv[e]	HC[f]	Senti[g]	Lexicon[h]
Twitter	[51-54]	[53-56]	[51-53, 56]	[51, 54, 55]	[51-54, 56]	-	-	-
OHC[i]	-	[50]	[50]	-	-	-	-	-

[a]Supervised machine learning; [b]Unsupervised machine learning; [c]Classical machine learning; [d]Modern machine learning; [e]Surveillance; [f]Health communication; [g]Sentiment analysis; [h]Lexicon-based methods; [i]Online Health Communities

3.6 Mental Health

Mental health problems are estimated to account for 13% of the global burden of disease, as measured in Disability Adjusted Life Years [95]. Using social media as a resource to understand mental health is a research area that has experienced substantial growth in recent years [96], given the burden of disease associated with mental health problems and the fact that social media provides ready access to first person reports of behaviour, thoughts, and feelings. Reviewed studies covered a range of mental health topics, including predicting depression diagnosis [8], assessing suicide risk [16, 18, 24, 74-76, 98, 99], and developing a better understanding of users' experiences of eating disorders [15], schizophrenia [59, 61], grief processes between gang-involved youth [58], relaxation [62], stress [63], pathological empathy [67, 72], and negative emotional effects associated with campus-based mass murders [64]. Related to this, a range of metrics have been used to characterize language use associated with specific mental health conditions, with *lexical diversity*, *readability scores*, *sentence complexity*, *negation*, *uncertainty,* and *degree of repetition,* all used during the review period [23, 26, 27, 60]. In novel work focussing on the relationship between clinical guidelines and actual treatments, Zhang et al. [71] created a catalogue of real-world treatments used — as opposed to merely discussed — by parents of children with autistic spectrum disorder, and then automatically identified their frequency of mention in two online autism forums.

With a view to improving how mental health forums are designed, one study applied textual cluster analysis to forums related to the conditions anxiety, depression, and post-traumatic stress disorder (PTSD) [19], showing that — consistent with current thinking regarding the relationship between PTSD and anxiety [97] — anxiety and PTSD forums shared more similarities to each other than to the depression forum. Related to this, another study found that different communities provided different degrees of emotional and informational support [20], with some communities (e.g., depression forums) focussed primarily on emotional support, and other communities (e.g. obsessive compulsive disorder forums) offering a greater proportion of informational support. Furthermore, the same study found that at the user level, the provision of social support was correlated with demonstrated linguistic accommodation, suggesting that those users who were able to "match" the linguistic culture of a particular community were likely to receive a greater volume of social support. Finally, a further study [100] involved the development of a classifier capable of identifying respectful uses of a mental-health related term (e.g. "I'm fuming. How dare a TV show portray folks suffering from mental health issues so unfairly") and less-respectful usage.

Of the thirty-one mental health-related papers reviewed (see Table 8), thirteen involved the use of Reddit data [15-27], ten used Twitter data [18, 24, 58-65], one used Instagram [18], three used Facebook [8, 18, 67], six used OHC data [70-75], and one used data derived from Weibo [76], with twenty-two of the papers utilising supervised machine learning methods [8, 16, 18, 20-22, 24, 25, 58-62, 65, 67, 70-76], and twelve papers utilising unsupervised machine learning [8, 15, 18-22, 27, 59, 60, 70, 72]. The majority of the papers reported on the use of classical machine learning approaches [8, 15, 16, 18-20, 22, 24, 25, 27, 58-62, 65, 67, 71, 73-76], with a minority using modern machine learning methods [18, 21, 22, 67, 70, 72]. Four of the mental health papers reviewed utilised primarily lexicon-based methods [17, 23, 63, 64].

3.7 Ethical Issues

Two types of ethics-related papers are discussed in this section: those that are focussed on empirical ethics (i.e. the empirical investigation of ethical beliefs and practices) [101, 102], and those that are focussed on ethical guideline development (i.e. the generation of theoretical frameworks and practical guidelines for conducting health-related NLP research with social media) [9, 103, 104]. Reviewed studies highlighted the need for both transparency in the development of algorithms and an ethical framework to guide the appropriate use of social media for computational public health research.

Focussing specifically on research ethics from the perspective of social media *users*, one study [102] pointed to a generally favourable view of the use of computational methods for public health research among social media users, provided that data was highly aggregated, and the goal of the work was of significant public health value (e.g. opioid abuse surveillance was acceptable in a public health context, but not when used for employment screening). However, among some users, concerns remained regarding the robustness of both the data and the research methods, due to the fact that the data was not representative of the general population, and was subject to impression management (i.e. many users did not tweet about stigmatising health problems [105]). Related to this work, one paper — a systematic review of attitudes towards the ethics of computational social media research [106] — found a range of different views on appropriate research ethics, depending on the particular research topic discussed, suggesting that a "blanket" approach to research ethics is currently not appropriate, and instead ethical deliberations ought to take into account the particular context of the research under review [106].

As noted by Vayena et al., [104], the research regulation infrastructure in most jurisdictions was developed in the period prior to social media, and hence is not well-equipped to manage the review of computational social media research. This point is reinforced by a qualitative study conducted with Research Ethics Committee (Institutional Review Board) members in the United Kingdom. This study outlines the challenges faced by ethics committees in the application of existing research ethics regulation to computational work and emphasises the need to protect research participants (i.e. social media users), even in the context of research using publicly available data [101].

Finally, practical guidelines have recently been developed to guide NLP research using social media data [103], with eight principles outlined, including the stipulation that as most social media based NLP research can be defined as human subjects research [107], ethical approval or exemption ought to be gained from an Institutional Review Board or Research Ethics Committee; that data ought to be de-identified for use in publications and presentations; and that caution ought to be exercised in linking data.

In recent years there has been a move away from the commonly held view that in social media research "anything goes", towards a more sophisticated perspective that acknowledges both the existence and importance of the ethical and regulatory issues involved in the application of NLP to social media for health research. Further, the provision of ethical guidelines developed specifically for NLP researchers — as described above, [103] — is a new and welcome development in the period since 2016.

Table 8 Summary of mental health-related papers

Data source	SML[a]	UML[b]	CML[c]	MML[d]	Surv[e]	HC[f]	Senti[g]	Lexicon[h]
Reddit	[16, 18, 20-22, 24, 25]	[15, 18-22, 27]	[15, 16, 18-20, 22, 24, 25, 27]	[21, 22]	–	–	[26]	[17], [23]
Twitter	[18, 58-62, 65]	[18, 59, 60]	[58-62, 65]	[18]	–	–	[63, 64, 24]	[63, 64]
Instagram	[18]	[18]	–	[18]	–	–	–	–
Facebook	[8, 18, 67]	[8, 18]	[8, 67]	[18, 67]	–	–	–	–
OHC[i]	[70-75]	[70, 72]	[71, 73-75]	[70, 72]	–	–	–	–
Weibo	[76]	–	[76]	–	–	–	–	–

[a]Supervised machine learning; [b]Unsupervised machine learning; [c]Classical machine learning; [d]Modern machine learning; [e]Surveillance; [f]Health communication; [g]Sentiment analysis; [h]Lexicon-based methods; iOnline Health Communities

4 Discussion and Conclusion

In this survey, we have presented recent advances in the application of NLP to social media to address public health research questions. We observed a substantial growth in the area of mental health and substance abuse research, and a continuing sustained interest in the use of social media for studying communicable diseases (particularly in the area of vaccine hesitancy). The widespread use of lexical resources developed in the psychology research communities — specifically, LIWC — is also notable, as is the relatively low frequency of "modern" (as opposed to "classical") machine learning approaches.

While predicting future trends is not a straightforward task, we tentatively suggest four directions in which current work is evolving. First, linking data — with appropriate consent — from the EHR and social media, both in the context of public health research and clinical care. Examples of this type of work in the research context already exist (e.g. [8]), and will likely be a focus of considerable research effort over the next few years.

Second, further utilisation of social media in public health surveillance. Currently, while advances have been made in

research using NLP and social media, substantial barriers still exist to implementing social media health surveillance in the context of public health practice. These barriers include costs (public health agencies are frequently underfunded), limited expertise in NLP, and difficulties in integrating social media analysis with existing surveillance methods and pipelines. However, even given these challenges, considerable strides have been made, particularly in the area of pharmacovigilance (e.g. the Food & Drug Administration Center for Drug Evaluation and Research).

Third, much social media research relies on the identification of appropriate keywords to construct a data sample suitable for the research question at hand. This keyword selection process has typically relied on intuition. However, recently there has been a move towards a more data-driven means of iteratively identifying and evaluating keywords (and their associated synonyms), with word embeddings and other empirical synonym discovery methods (e.g. [108]). This shift towards a more principled method of selecting keywords for data sampling is to be welcomed.

Fourth, while we believe that Twitter will remain a valuable (and popular) data source for NLP research, we suspect that Reddit will become increasingly popular as a research resource, partly due to its "research-friendly" terms and conditions and its increasing user base. Related to this, the dynamism of the social media ecosystem should not be underestimated, with new services (e.g. TikTok) attracting users — especially new adolescent users — away from existing services. Given this rapidly changing social media environment, there is little reason to believe that currently popular social media platforms will maintain their current level of popularity.

Acknowledgements

This work was partially supported by the National Institute on Drug Abuse of the United States National Institutes of Health under award number R21DA043775. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

References

1. Fox S. The social life of health information; Available from: http://www.pewresearch.org/fact-tank/2014/01/15/the-social-life-of-health-information/.
2. Paul M, Dredze M. Social Monitoring for Public Health. Marchionini G, editor. Morgan Claypool; 2017.
3. Guntuku S, Yaden D, Kern M, Ungar L, Eichstaedt J. Detecting depression and mental illness on social media: an integrative review. Curr Opin Behav Sci 2017;18:43–9.
4. Trautmann S, Rehm J, Wittchen HU. The economic costs of mental disorders: Do our societies react appropriately to the burden of mental disorders? EMBO Rep 2016 Sep;17(9):1245-9.
5. Tangherlini TR, Roychowdhury V, Glenn B, Crespi CM, Bandari R, Wadia A, et al. "Mommy Blogs" and the vaccination exemption narrative: results from a machine-learning approach for story aggregation on parenting social media sites. JMIR Public Health Surveill 2016 Nov;2(2):e166.
6. Allem JP, Dharmapuri L, Unger JB, Cruz TB. Characterizing JUUL-related posts on Twitter. Drug Alcohol Depend 2018 09;190:1–5.
7. Charles-Smith L, Reynolds T, Cameron M, Conway M, Lau E, Olsen J, et al. Using social media for actionable disease surveillance and outbreak management: a systematic literature review. PLoS One 2015;10(10):e0139701.
8. Eichstaedt JC, Smith RJ, Merchant RM, Ungar LH, Crutchley P, Preoţiuc-Pietro D, et al. Facebook language predicts depression in medical records. Proc Natl Acad Sci U S A 2018 Oct;115(44):11203–8.
9. Vayena E, Blasimme A, Cohen IG. Machine learning in medicine: Addressing ethical challenges. PLoS Med 2018 Nov;15(11):e1002689.
10. Park A, Conway M. Tracking health related discussions on Reddit for public health applications. AMIA Annu Symp Proc 2017;2017:1362–71.
11. Meacham M, Paul M, Ramo D. Understanding emerging forms of cannabis use through an online cannabis community: An analysis of relative post volume and subjective highness ratings. Drug Alcohol Depend 2018 07;188:364–9.
12. Zhan Y, Liu R, Li Q, Leischow S, Zeng D. Identifying topics for e-cigarette user- generated contents: a case study from multiple social media platforms. J Med Internet Res 2017 01;19(1):e24.
13. Chen A, Zhu SH, Conway M. What online communities can tell us about electronic cigarettes and hookah use: a study using text mining and visualization techniques. J Med Internet Res 2015 Sep;17(9):e220.
14. Nobles A, Dreisbach C, Kelm-Malpass J, Barnes L. "Is This an STD? Please Help!": Online Information Seeking for Sexually Transmitted Diseases on Reddit. In: Proceedings of the Twelfth International Conference on Web and Social Media; 2018. p. 660–3.
15. Moessner M, Feldhege J, Wolf M, Bauer S. Analyzing big data in social media: text and network analyses of an eating disorder forum. Int J Eat Disord 2018 Jul;51(7):656–67.
16. Aladağ A, Muderrisoglu S, Akbas N, Zahmacioglu O, Bingol H. Detecting suicidal Ideation on forums: proof-of-concept study. J Med Internet Res.2018 Jun;20(6):e215.
17. Park A, Conway M. Harnessing Reddit to understand the written-communication challenges experienced by individuals with mental health disorders: analysis of texts from mental health communities. J Med Internet Res 2018 Apr;20(4):e121.
18. Coppersmith G, Leary R, Crutchley P, Fine A. Natural language processing of social media as screening for suicide risk. Biomed Inform Insights 2018;10:1178222618792860.
19. Park A, Conway M, Chen A. Examining thematic similarity, difference, and membership in three online mental health communities from Reddit: a text mining and visualization approach. Comput Human Behav 2018 Jan;78:98–112.
20. Sharma E, De Choudhury M. Mental health support and its relationship to linguistic accommodation in online communities. In: Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems. CHI '18. New York, NY, USA: ACM; 2018. p. 641:1–641:13. Available from: http://doi.acm.org/10.1145/3173574.3174215.
21. Ive J, Gkotsis G, Dutta R, Stewart R, Velupillai S. Hierarchical neural model with attention mechanisms for the classification of social media text related to mental health. In: Proceedings of the Fifth Workshop on Computational Linguistics and Clinical Psychology: From Keyboard to Clinic. Association for Computational Linguistics; 2018. p. 69–77. Available from: http://aclweb.org/anthology/W18-0607.
22. Gkotsis G, Oellrich A, Velupillai S, Liakata M, Hubbard T, Dobson R, et al. Characterisation of mental health conditions in social media using Informed Deep Learning. Sci Rep 2017 03;7:45141.
23. Park A, Conway M. Longitudinal changes in psychological states in online health community members: understanding the long-term effects of participating in an online depression community. J Med Internet Res 2017 03;19(3):e71.
24. Kavuluru R, Williams AG, Ramos-Morales M, Haye L, Holaday T, Cerel J. Classification of helpful comments on online suicide watch forums. ACM BCB 2016 Oct;2016:32– 40.
25. De Choudhury M, Kiciman E, Dredze M, Coppersmith G, Kumar M. Discovering shifts to suicidal ideation from mental health content in social media. Proc SIGCHI Conf Hum Factor Comput Syst 2016 May;2016:2098–110.
26. Gkotsis G, Oellrich A, Hubbard T, Dobson R, Liakata M, Velupillai S, et al. The language of mental health problems in social media. In: Proceedings of the Third Workshop on Computational Lingusitics and Clinical Psychology. Association for Computational Linguistics; 2016. p. 63–73. Available from: http://aclweb.org/anthology/W16-0307.
27. Kumar M, Dredze M, Coppersmith G, De Choudhury M. Detecting changes in suicide content manifested in social media following celebrity suicides. HT ACM Conf Hypertext Soc Media

2015 Sep;2015:85–94.
28. Mitra T, Counts S, Pennebaker J. Understanding anti-vaccination attitudes in social media. In: Proceedings of the Tenth International Conference on Web and Social Media (ICWSM 2016); 2016. p. 269-78.
29. Surian D, Nguyen DQ, Kennedy G, Johnson M, Coiera E, Dunn AG. Characterizing Twitter discussions about HPV vaccines using topic modeling and community detection. J Med Internet Res 2016 08;18(8):e232.
30. Massey P, Leader A, Yom-Tov E, Budenz A, Fisher K, Klassen A. Applying multiple data collection tools to quantify human papillomavirus vaccine communication on Twitter. J Med Internet Res 2016 12;18(12):e318.
31. Wakamiya S, Kawai Y, Aramaki E. Twitter-based influenza detection after flu peak via tweets with indirect information: text mining study. JMIR Public Health Surveill 2018 Sep;4(3):e65.
32. Al-Garadi MA, Khan MS, Varathan KD, Mujtaba G, Al-Kabsi AM. Using online social networks to track a pandemic: A systematic review. J Biomed Inform 2016 08;62:1–11.
33. Mowery J. Twitter Influenza Surveillance: Quantifying Seasonal Misdiagnosis Pat- terns and their Impact on Surveillance Estimates. Online J Public Health Inform 2016;8(3):e198.
34. Zhang L, Hall M, Bastola D. Utilizing Twitter data for analysis of chemotherapy. Int J Med Inform 2018 Dec;120:92–100.
35. Mackey T, Kalyanam J, Klugman J, Kuzmenko E, Gupta R. Solution to detect, classify, and report illicit online marketing and sales of controlled substances via Twitter: using machine learning and web forensics to combat digital opioid access. J Med Internet Res 2018 04;20(4):e10029.
36. Curtis B, Giorgi S, Buffone AEK, Ungar LH, Ashford RD, Hemmons J, et al. Can Twitter be used to predict county excessive alcohol consumption rates? PLoS One 2018;13(4):e0194290.
37. Simpson SS, Adams N, Brugman CM, Conners TJ. Detecting novel and emerging drug terms using natural language processing: a social media corpus study. JMIR Public Health Surveill 2018 Jan;4(1):e2.
38. Ayers J, Dredze M, Leas E, Caputi T, Allem JP, Cohen J. Next generation media monitoring: Global coverage of electronic nicotine delivery systems (electronic cigarettes) on Bing, Google and Twitter, 2013-2018. PLoS One 2018;13(11):e0205822.
39. Mackey TK, Kalyanam J, Katsuki T, Lanckriet G. Twitter-based detection of illegal online sale of prescription opioid. Am J Public Health 2017 12;107(12):1910–1915.
40. Huang T, Elghafari A, Relia K, Chunara R. High-resolution temporal representations of alcohol and tobacco behaviors from social media data. Proc ACM Hum Comput Interact 2017 Nov;1(CSCW).
41. Glowacki EM, Glowacki JB, Wilcox GB. A text-mining analysis of the public's reactions to the opioid crisis. Subst Abus 2017 Jul;p. 1–5.
42. Mackey TK, Kalyanam J. Detection of illicit online sales of fentanyls via Twitter. F1000Res 2017;6:1937.
43. Lazard A, Saffer A, Wilcox G, Chung AD, Mackert M, Bernhardt J. E-cigarette social media messages: a text mining analysis of marketing and consumer conversations on Twitter. JMIR Public Health Surveill 2016 Dec;2(2):e171.
44. Kolliakou A, Ball M, Derczynski L, Chandran D, Gkotsis G, Deluca P, et al. Novel psychoactive substances: An investigation of temporal trends in social media and electronic health records. Eur Psychiatry 2016 10;38:15–21.
45. Kostygina G, Tran H, Shi Y, Kim Y, Emery S. "Sweeter Than a Swisher": amount and themes of little cigar and cigarillo content on Twitter. Tob Control 2016 10;25(Suppl 1):i75–i82.
46. Daniulaityte R, Chen L, Lamy FR, Carlson RG, Thirunarayan K, Sheth A. "When 'Bad' is 'Good'": Identifying Personal Communication and Sentiment in Drug-Related Tweets. JMIR Public Health Surveill 2016 Oct;2(2):e162.
47. Kavuluru R, Sabbir AKM. Toward automated e-cigarette surveillance: Spotting e- cigarette proponents on Twitter. J Biomed Inform 2016 06;61:19–26.
48. Alvaro N, Conway M, Doan S, Lofi C, Overington J, Collier N. Crowdsourcing Twitter annotations to identify first-hand experiences of prescription drug use. J Biomed Inform 2015 Dec;58:280–7.
49. Kim A, Hopper T, Simpson S, Nonnemaker J, Lieberman A, Hansen H, et al. Using Twitter data to gain insights into e-cigarette marketing and locations of use: An infoveillance study. J Med Internet Res 2015 Nov;17(11):e251.
50. Park S, Hong S. Identification of primary medication concerns regarding thyroid hormone replacement therapy From online patient medication reviews: text mining of social network data. J Med Internet Res 2018 Oct;20(10):e11085.
51. Sarker A, Belousov M, Friedrichs J, Hakala K, Kiritchenko S, Mehryary F, et al. Data and systems for medication-related text classification and concept normalization from Twitter: insights from the Social Media Mining for Health (SMM4H)-2017 shared task. J Am Med Inform Assoc 2018 Oct;25(10):1274–83.
52. Bollegala D, Maskell S, Sloane R, Hajne J, Pirmohamed M. Causality patterns for detecting adverse drug reactions from social media: text mining approach. JMIR Public Health Surveill 2018 May;4(2):e51.
53. Kagashe I, Yan Z, Suheryani I. Enhancing seasonal influenza surveillance: topic analysis of widely used medicinal drugs using Twitter data. J Med Internet Res 2017 Sep;19(9):e315.
54. Cocos A, Fiks A, Masino A. Deep learning for pharmacovigilance: recurrent neural network architectures for labeling adverse drug reactions in Twitter posts. J Am Med Inform Assoc 2017 Jul;24(4):813–21.
55. Sarker A, Gonzalez G. A corpus for mining drug-related knowledge from Twitter chatter: Language models and their utilities. Data Brief 2017 Feb;10:122–131.
56. Sarker A, O'Connor K, Ginn R, Scotch M, Smith K, Malone D, et al. Social media mining for toxicovigilance: automatic monitoring of prescription medication abuse from Twitter. Drug Saf 2016 Mar;39(3):231–40.
57. Young SD, Mercer N, Weiss RE, Torrone EA, Aral SO. Using social media as a tool to predict syphilis. Prev Med 2018 Apr;109:58–61.
58. Patton DU, MacBeth J, Schoenebeck S, Shear K, McKeown K. Accommodating grief on Twitter: an analysis of expressions of grief among gang involved youth on Twitter using qualitative analysis and natural language processing. Biomed Inform Insights 2018;10:1178222618763155.
59. Ernala S, Labetoulle T, Bane F, Bimbaum M, Rizvi A, Kane J, et al. Characterizing Audience Engagement and Assessing Its Impact on Social Media Disclosures of Mental Illnesses. In: Proceedings of the Twelfth International Conference on Web and Social Media; 2018. p. 62–71.
60. Guntuku SC, Ramsay JR, Merchant RM, Ungar LH. Language of ADHD in adults on social media. J Atten Disord 2017 Nov:1087054717738083.
61. Birnbaum M, Ernala S, Rizvi A, De Choudhury M, Kane J. A collaborative approach to identifying social media markers of schizophrenia by employing machine learning and clinical appraisals. J Med Internet Res 2017 Aug;19(8):e289.
62. Doan S, Ritchart A, Perry N, Chaparro J, Conway M. How do You relax When You're stressed? a content analysis and infodemiology study of stress-related tweets. JMIR Public Health Surveill 2017 Jun;3(2):e35.
63. Loveys K, Crutchley P, Wyatt E, Coppersmith G. Small but mighty: affective micropatterns for quantifying mental health from social media language. In: Proceedings of the Fourth Workshop on Computational Linguistics and Clinical Psychology — From Linguistic Signal to Clinical Reality. Association for Computational Linguistics; 2017. p. 85–95. Available from: http://aclweb.org/anthology/W17-3110.
64. Jones N, Wojcik S, Sweeting J, Silver RC. Tweeting negative emotion: an investigation of Twitter data in the aftermath of violence on college campuses. Psychol Methods 2016 12;21(4):526-41.
65. Mowery D, Park A, Bryan C, Conway M. Towards Automatically Classifying Depressive Symptoms from Twitter Data for Population Health. In: Proceedings of the Workshop on Computational Modeling of People's Opinions, Personality, and Emotions in So- cial Media (PEOPLES). The COLING 2016 Organizing Committee; 2016. p. 182–191. Available from: http://aclweb.org/anthology/W16-4320.
66. Faasse, K., Chatman, C., and Martin, L. A comparison of language use in pro- and anti- vaccination comments in response to a high profile Facebook post. Vaccine. 2016 11;34(47):5808–5814.
67. Abdul-Mageed M, Buffone A, Peng H, Eichstaedt J, Ungar L. Recognizing pathogenic empathy in social media. In: Proceedings of the Eleventh International Conference on Web and Social Media; 2017. p. 448–51.
68. Zhang S, Grave E, Sklar E, Elhadad N. Longitudinal analysis of discussion topics in an online breast cancer community using convolutional neural networks. J Biomed Inform 2017 05;69:1–9.
69. Zhang S, Qiu L, Chen F, Zhang W, Yu Y, Elhadad N. "We make choices we think are going to save us": Debate and stance identification for online breast cancer CAM discussions. Proc Int World Wide Web Conf 2017 Apr;2017:1073–81.

70. Khanpour H, Caragea C. Fine-grained emotion detection in health-related online posts. In: Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing. Association for Computational Linguistics; 2018. p. 1160–6. Available from: http://aclweb.org/anthology/D18-1147.
71. Zhang S, Kang T, Qiu L, Zhang W, Yu Y, Elhadad N. Cataloguing treatments dis- cussed and used in online autism communities. Proc Int World Wide Web Conf 2017 Apr;2017:123–31.
72. Khanpour H, Caragea C, Biyani P. Identifying empathetic messages in online health ommunities. In: Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 2: Short Papers). Asian Federation of Natural Language Processing; 2017. p. 246–51. Available from: http://aclweb.org/anthology/ I17-2042.
73. Franco-Penya H, Mamani Sanchez L. Text-based experiments for predicting mental health emergencies in online web forum posts. In: Proceedings of the Third Work- shop on Computational Lingusitics and Clinical Psychology. Association for Computational Linguistics; 2016. p. 193–7. Available from: http://aclweb.org/anthology/W16-0327.
74. Asgari E, Nasiriany S, Mofrad M. Text Analysis and Automatic Triage of Posts in a Mental Health Forum. In: Proceedings of the Third Workshop on Computational Lingusitics and Clinical Psychology. Association for Computational Linguistics; 2016. p. 153–7. Available from: http://aclweb.org/anthology/W16-0318.
75. Cohan A, Young S, Goharian N. Triaging mental health forum posts. In: Proceedings of the Third Workshop on Computational Lingusitics and Clinical Psychology. Association for Computational Linguistics; 2016. p. 143–7. Available from: http://aclweb.org/ anthology/W16-0316.
76. Cheng Q, Li TM, Kwok CL, Zhu T, Yip PS. Assessing suicide risk and emotional distress in Chinese social media: a text mining and machine learning study. J Med Internet Res 2017 Jul;19(7):e243.
77. Guo H, Na X, Hou L, Li J. Classifying Chinese Questions Related to Health Care Posted by Consumers Via the Internet. J Med Internet Res 2017 Jun;19(6):e220.
78. Boden M. Mind as Machine: A History of Cognitive Science. OUP; 2006.
79. Tausczik Y, Pennebaker J. The psychological meaning of words: LIWC and computerized text analysis methods. J Lang Soc Psychol 2010;29(1):24– 54.
80. Brownstein CJ, Sand F. HealthMap: the development of automated real-time internet surveillance for epidemic intelligence. Euro Surveill 2007 Nov;12(11):E071129.5.
81. Collier N, Doan S, Kawazoe A, Matsuda-Goodwin R, Conway M, Tateno Y, et al. BioCaster: detecting public health rumors with a Web-based text mining system. Bioinformatics 2008 Dec;24(24):2940–1.
82. Eysenbach G. Infodemiology and infoveillance: framework for an emerging set of public health informatics methods to analyze search, communication and publication behavior on the Internet. J Med Internet Res 2009 Mar;11(1):e11.
83. Brownstein JS, Freifeld CC, Madoff LC. Digital disease detection–harnessing the Web for public health surveillance. N Engl J Med 2009 May;360(21):2153–5, 2157.
84. Salathé M, Bengtsson L, Bodnar TJ, Brewer DD, Brownstein JS, Buckee C, et al. Digital epidemiology. PLoS Comput Biol 2012;8(7):e1002616.
85. Singh T, Arrazola RA, Corey CG, Husten CG, Neff LJ, Homa DM, et al. Tobacco use among middle and high school students–United States, 2011-2015. MMWR Morb Mortal Wkly Rep 2016 Apr;65(14):361–7.
86. Grana R, Benowitz N, Glantz SA. E-cigarettes: a scientific review. Circulation 2014 May;129(19):1972–86.
87. McNeill A, Brose L, Calder R, Hitchman S. E-cigarettes: An Evidence Update - Report Commissioned by Public Health England. Public Health England; 2015.
88. Polosa R. E-cigarettes: Public Health England's evidence based confusion? Lancet 2015 Sep;386(10000):1237–8.
89. Pacula RL, Powell D, Heaton P, Sevigny EL. Assessing the effects of medical marijuana laws on marijuana use: the devil is in the details. J Policy Anal Manage 2015;34(1):7– 31.
90. Hasin DS, Shmulewitz D, Sarvet AL. Time trends in US cannabis use and cannabis use disorders overall and by sociodemographic subgroups: a narrative review and new findings. Am J Drug Alcohol Abuse 2019 Mar 14:1–21.
91. GBD 2016 Alcohol Collaborators. Alcohol use and burden for 195 countries and territories, 1990-2016: a systematic analysis for the Global Burden of Disease Study 2016. Lancet 2018 09;392(10152):1015–35.
92. Paul M, Sarker A, Brownstein J, Nikfarjam A, Scotch M, Smith K, et al. Social media mining for public health monitoring and surveillance. In: Proceedings of the Pacific Symposium on Biocomputing 2016; 2016. p. 468–79.
93. Bigeard E, Grabar N, Thiessard F. Detection and analysis of drug misuses. a study based on social media messages. Front Pharmacol 2018;9:791.
94. Manchikanti L, Helm S, Fellows B, Janata J, Pampati V, Grider J, et al. Opioid epidemic in the United States. Pain Physician 2012 Jul;15(3 Suppl):ES9–38.
95. Vigo D, Thornicroft G, Atun R. Estimating the true global burden of mental illness. Lancet Psychiatry 2016 Feb;3(2):171–8.
96. Conway M, O'Connor D. Social media, big data, and mental health: current advances and ethical implications. Curr Opin Psychol 2016 Jun;9:77–82.
97. Ginzburg K, Ein-Dor T, Solomon Z. Comorbidity of posttraumatic stress disorder, anxiety and depression: a 20-year longitudinal study of war veterans. J Affect Disord 2010 Jun;123(1-3):249–57.
98. Nock M, editor. The Oxford Handbook of Suicide and Self-Injury. OUP; 2014.
99. Bryan C, Butner J, Sinclair S, Bryan AB, Hesse C, Rose A. Predictors of emerging suicide death among military personnel on social media networks. Suicide Life Threat Behav 2018 08;48(4):413–30.
100. Hwang JD, Hollingshead K. Crazy mad nutters: the language of mental health. In: Proceedings of the Third Workshop on Computational Lingusitics and Clinical Psychology. Association for Computational Linguistics; 2016. p. 52–62. Available from: http://aclweb.org/anthology/W16-0306.
101. Hibbin RA, Samuel G, Derrick GE. From "a fair game" to "a form of covert research": research ethics committee members' differing notions of consent and potential risk to participants within social media research. J Empir Res Hum Res Ethics 2018 Apr;13(2):149–159.
102. Mikal J, Hurst S, Conway M. Ethical issues in using Twitter for population-level depression monitoring: a qualitative study. BMC Med Ethics 2016 Apr;17:22.
103. Benton A, Coppersmith G, Dredze M. Ethical research protocols for social media health research. In: Proceedings of the First ACL Workshop on Ethics in Natural Language Processing. Association for Computational Linguistics; 2017. p. 94–102. Available from: http://aclweb.org/anthology/W17-1612.
104. Vayena E, Salathé M, Madoff LC, Brownstein JS. Ethical challenges of big data in public health. PLoS Comput Biol 2015 Feb;11(2):e1003904.
105. Goffman E. Stigma: Notes on the Management of Spoiled Identity. A Spectrum book. Englewood Cliffs, N.J.: Prentice-Hall; 1963.
106. Golder S, Ahmed S, Norman G, Booth A. Attitudes toward the ethics of research using social media: a systematic review. J Med Internet Res 2017 Jun;19(6):e195.
107. O'Connor D. The apomediated world: regulating research when social media has changed research. Journal of Law, Medicine, and Ethics 2013;41(2):470–83.
108. Adams N, Artigiani E, Wish E. Choosing your platform for social media drug research and improving your keyword filter list. Journal of Drug Issues 2019;1(16).

Correspondence to:
Mike Conway
E-mail: mike.conway@utah.edu