

Mining Social Media Data for Biomedical Signals and Health-Related Behavior

Rion Brattig Correia,^{1,2,3} Ian B. Wood,² Johan Bollen,²
and Luis M. Rocha^{1,2}

¹Instituto Gulbenkian de Ciéncia, 2780-156 Oeiras, Portugal

²Center for Social and Biomedical Complexity, Luddy School of Informatics, Computing & Engineering, Indiana University, Bloomington, Indiana 47408, USA; email: rocha@indiana.edu

³CAPES Foundation, Ministry of Education of Brazil, 70040 Braslia DF, Brazil

Ann. Rev. Biomed. Data Sci. 2020. 3:433–58

First published as a Review in Advance on
May 4, 2020

The *Annual Review of Biomedical Data Science* is
online at biodatasci.annualreviews.org

<https://doi.org/10.1146/annurev-biodatasci-030320-040844>

Copyright © 2020 by Annual Reviews.
All rights reserved

ANNUAL REVIEWS **CONNECT**

www.annualreviews.org

- Download figures
- Navigate cited references
- Keyword search
- Explore related articles
- Share via email or social media

Keywords

social media, healthcare, pharmacovigilance, sentiment analysis, biomedicine

Abstract

Social media data have been increasingly used to study biomedical and health-related phenomena. From cohort-level discussions of a condition to population-level analyses of sentiment, social media have provided scientists with unprecedented amounts of data to study human behavior associated with a variety of health conditions and medical treatments. Here we review recent work in mining social media for biomedical, epidemiological, and social phenomena information relevant to the multilevel complexity of human health. We pay particular attention to topics where social media data analysis has shown the most progress, including pharmacovigilance and sentiment analysis, especially for mental health. We also discuss a variety of innovative uses of social media data for health-related applications as well as important limitations of social media data access and use.

1. INTRODUCTION

Humanity has crossed an important threshold in its ability to construct quantitative, large-scale characterizations of the networks of information exchanges and social interactions in human societies. Due to the widespread digitization of behavioral and medical data, the advent of social media, and the Internet's infrastructure of large-scale knowledge storage and distribution, there has been a breakthrough in our ability to characterize human social interactions, behavioral patterns, and cognitive processes and their relationships with biomedicine and healthcare. For instance, electronic health records (EHRs) of entire cities can yield valuable insights on gender and age disparities in healthcare (1), and communication patterns on Twitter and Instagram can help us detect the spread of flu pandemics (2), warning signals of drug interactions (3), and depression (4).

Data science, together with artificial intelligence and complex networks and systems theory, has already enabled exciting developments in the social sciences, including the novel fields of computational social science and digital epidemiology (5, 6). Using social media and online data, researchers in these interdisciplinary fields are tackling human behavior and society in a large-scale quantitative manner not previously possible to study, for example, social protests (7, 8), fake news spread (9), and stock market prediction (10).

This approach also shows great promise in monitoring human health and disease given the newfound ability to measure the behavior of very large populations using individual self-reports (11). Indeed, population-level observation tools allow us to study collective human behavior (12–14) and, given the ability to obtain large amounts of real-world behavioral data, are expected to speed translational research in transformative ways (15) by monitoring of individual and population health (15–19). This promise has been substantiated by many recent studies. Google searches have been shown to correlate with dengue spread in tropical zones, for example (20). Although the accuracy of Google trends data alone has been problematic for epidemic flu modeling (21), such data add value in combination with other health data (22). Several studies have also shown that social media analysis is useful for: tracking and predicting disease outbreaks such as influenza (11, 23), cholera (24), Zika (25), and HIV (26); playing an important role in pharmacovigilance (3, 27–31); and measuring public sentiment and other signals associated with public health issues such as depression (4, 32–34), human reproduction (35), vaccination rates (36, 37), and mental disorder stigma (13).

Social media data provide an increasingly detailed large-scale record of the behavior of a considerable fraction (about one-seventh) of the world's population. Since 2017, 330 million people monthly have been active users of Twitter, making it one of the most populated global social networking platforms (38). Instagram currently has more than one billion monthly active users (39). It used to be the preferred social network among teens and young adults (ages 12–24), but since 2016 Instagram was surpassed by Snapchat in this demographic (40). Facebook, however, still has the overall majority of active users, with 2.45 billion users monthly (41) and a total of 2.8 billion across all of the company's core products: Facebook, WhatsApp, Instagram, and Messenger (42).

Biomedical and public health researchers now have the ability to directly measure human behavior on social media, a promise emphasized by the National Institutes of Health, which consider this type of big data very relevant for biomedical research (43, 44). By "social media" we mean any user-generated content, including posts to sites such as Twitter and Facebook but also comments on disease-specific or health-related sites, forums, or chats. Most social media sites have been shown to be relevant for biomedical studies, including Twitter (15), Facebook (12), Flickr (17), Instagram (3), Reddit (45–47), and even YouTube (48, 49). Used together with other sources of data such as web search, mobility data, scientific publications, EHRs, and genome-wide studies, social media data help researchers build population- and individual-level observation tools that can speed translational research in fundamentally new ways.

Leveraging these kinds of data constitutes a novel opportunity to improve personalization in the face of multilevel complexity in human disease (50). For instance, new patient-stratification principles and unknown disease correlations and comorbidities can now be revealed (51). Moreover, social media data allow for a more direct measurement of the perspective of patients on disease, which is often different from that of physicians. Social media can help both patients and practitioners to understand and reduce this disconnect (52), which is known to hinder treatment adherence (53).

In summary, analysis of social media data enables more accurate “microscopes” of individual human behavior and decision making, as well as “macroscopes” for collective phenomena (5, 54). These micro- and macro-level observation tools can go beyond a descriptive understanding of biomedical phenomena in human populations by enabling quantitative measurement and prediction of various processes, as reviewed below. The ability to study humans as their own model organism is now a more reasonable prospect than ever before. Here we review recent work pertaining to the mining of social media for health-related information, that is, biomedical, epidemiological, or any social phenomena data of relevance to the multilevel complexity of human health (55). The review is structured as follows: Section 2 covers the use of social media for pharmacovigilance, including adverse drug reactions (ADRs) and drug–drug interactions (DDIs); Section 3 addresses the use of sentiment analysis tools to characterize individual and population behavior, especially mental health; Section 4 looks at the analysis of social media data for a wide variety of health-related applications; Section 5 considers limitations of the use of social media data; and Section 6 contains a conclusion and final remarks.

2. PHARMACOVIGILANCE

It is estimated that every year the United States alone spends up to \$30.1 billion due to ADRs, with each individual reaction costing on average \$2,262 (56). More than 30% of ADRs are caused by DDIs that can occur when patients take two or more drugs concurrently (polypharmacy). The DDI phenomenon is also a worldwide threat to public health (1, 57), especially with increased polypharmacy in aging populations.

Most ADR and DDI surveillance is still conducted by analysis of physician reports to regulatory agencies and by mining databases of those reports, such as the FDA (US Food and Drug Administration) Adverse Event Reporting System (FAERS) (58). However, clinical surveillance suffers from under reporting (59), which can be caused by clinicians failing to note adverse events or downgrading the severity of patients’ symptoms (60). For example, it has been well documented that depression and pain are underassessed by clinicians and underreported by patients, and therefore are under- or inappropriately managed, especially in specific cohorts such as athletes (61, 62). Even when clinicians are specifically trained or required to use screening tools for ADR, in practice these are done in a reactionary fashion at the time of a healthcare visit (59).

Such problems in reporting can be improved using new ways to study ADR and DDI phenomena provided by recently available large-scale online data about human behavior. Given the number of users, social media data are likely to allow for automated proactive identification of issues as they develop rather than after they occur and potentially become severe. Thus, analysis of social media data can identify underreported pathology associated with ADRs and further contribute to improvements in population health (Figure 1 shows a sample of social media posts containing drug and symptom mentions). For instance, it has been shown that the combination of clinical FAERS reports and Internet search logs can improve the detection accuracy of ADRs by 19% (63), that discussions on Twitter related to glucocorticoid therapy reveal that insomnia and weight gain are more common adverse events than are reported in the UK regulator’s ADR database, and that more serious side effects are comparatively less discussed (52). Another study has

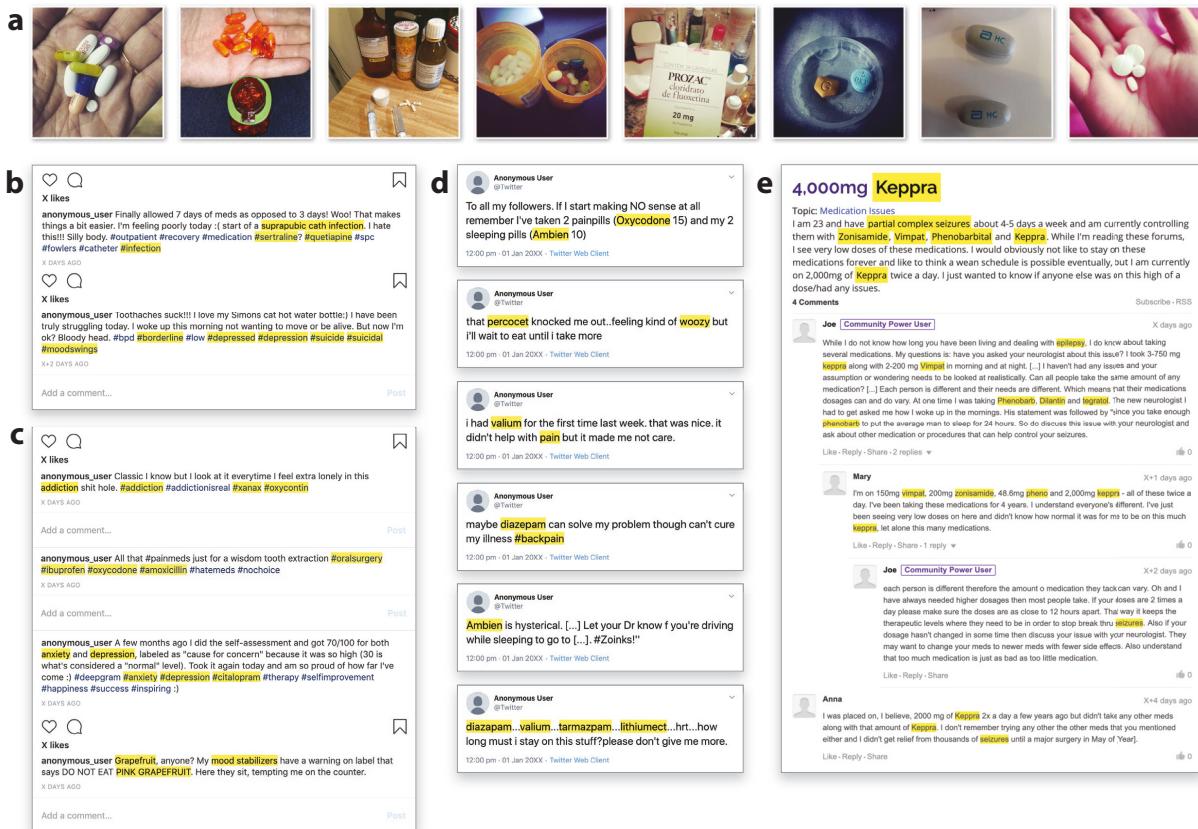


Figure 1

Selected sample of social media posts depicting known drug and symptom mentions. (a) Instagram photos depicting a variety of drugs. (b, c) Captions of Instagram posts. The two captions in panel b were posted two days apart by the same user, showing in the second post a possible side effect from a drug administration mentioned in the first post. (d) Twitter posts containing drugs known to be abused. (e) Epilepsy Foundation forum post and comments from users asking questions and sharing experiences over drug dosage (Keppra). For all examples, usernames, number of likes, and dates were omitted for privacy, and some content was modified for clarity and to maintain user anonymity. Terms of pharmacovigilance interest, including drug names, natural products, and symptoms, are highlighted in yellow using dictionaries developed for this problem (3, 65).

compared patient reports of ADRs on social media (various discussion forums on health-related websites) with those of clinicians on EHRs for the case of the two drugs aspirin and atorvastatin (64). This study found that the most frequently reported ADR in EHRs matched the patients' most frequently expressed concerns on social media. However, several less frequently reported reactions in EHRs were more prevalent on social media, with aspirin-induced hypoglycemia being discussed in social media only. The observed discrepancies and the increased accuracy and completeness of social media reports relative to those from regulator databases and EHRs have revealed that physicians and patients have different priorities (53). This suggests that social media data may be used to provide a more complete measurement of impact on quality of life (52), making it a useful complement to physician reporting.

The use of social media for pharmacovigilance is recent, but it has been receiving increasing attention in the last few years. A review paper in 2015 found only 24 studies on the topic,

almost evenly divided between manual and automated methods, and concluded that social media was likely useful for postmarketing drug surveillance (66). Another review paper from 2015 (29) found 22 studies and concluded that the utility of social media data analysis for biomedicine was hindered by the difficulty of comparing methods due to the scarcity of publicly available annotated data. This has led to a shared task workshop and the session Social Media Mining for Public Health Monitoring and Surveillance at the 2016 meeting of the Pacific Symposium on Biocomputing (15). Shared tasks have involved the automatic classification of posts mentioning ADRs, the extraction of related terms, and the normalization of standardized ADR lexicons (31). The conference session also attracted studies of social media data for a variety of health-related topics, including: tracking emotion (see Section 3) to detect disease outbreaks (67); pharmacovigilance, including dietary supplement safety (18) and ADR and DDI reports (3); and using machine learning (ML) to predict healthy behavior, such as diet success, using publicly shared fitness data from MyFitnessPal (68), smoking cessation from Twitter data (69), and overall well-being from volunteers' Facebook data (70). Since this initial event, the shared task and workshop, currently named Social Media Mining for Health Applications, has been held annually and serves to bring together researchers interested in automatic methods for the collection, extraction, representation, analysis, and validation of social media data for health informatics (71, 72).

Before the community was able to analyze well-known social media sites such as Twitter and Facebook, most pharmacovigilance work on mining ADRs from social media had been focused on social interactions in specialized health forums and message boards (73–79). Schatz's group was one of the first to pursue this angle by using network visualization, natural language processing (NLP), and sentiment analysis (see Section 3) to provide a qualitative ADR analysis of user comments on Yahoo Health Groups. They have shown that it is possible to visualize and monitor drugs and their ADRs in postmarketing (73), as well as to track patient sentiment regarding particular drugs over time (80).

Around the same time, Gonzalez's group created an ADR-focused lexicon and manually annotated a corpus of comments in support of a rule-based, lexical matching system developed to analyze user comments on DailyStrength (<https://www.dailystrength.org/>; a health-focused site where users discuss personal experiences with drugs) and demonstrated that comments contain useful drug safety information (74). Later the group used association rule mining to automatically extract ADRs from user comments in the same platform (75), other supervised classifiers to predict if individual comments contained ADRs, and a probabilistic model to infer if the DailyStrength footprint of such predicted ADRs for a given drug were likely to indicate a public health red flag (79).

Subsequently, Benton et al. (76) used co-occurrence statistics of drug–adverse effect pairs present in breast cancer message boards and compared them to drug labels of four different drugs. They found that 75–80% of these ADRs were documented on drug labels, while the rest were previously unidentified ADRs for the same drugs. Casting the extraction of (unreported) drug–event pairs in ADRs as a sequence labeling problem, Sampathkumar et al. (77) used a hidden Markov model on patient feedback data from Medications.com that had been automatically annotated using dictionaries of drug names, side effects, and interaction terms.

Several text mining and ML pipelines, as well as annotated corpora and lexica, were quickly developed for extraction and prediction of ADRs from various health forums and message boards. C. Yang et al. (81) used association mining and proportional reporting ratios to show that ADRs can be extracted from MedHelp (<https://www.medhelp.org/>) user comments; this study was conducted for a small set of five known ADRs (via FDA alerts) involving ten drugs. For the same platform, M. Yang et al. (82) used semisupervised text classification to filter comments likely to contain ADRs in support of an early warning pharmacovigilance system that they tested

successfully, albeit with only three drugs associated with more than 500 discussion threads. Yates & Goharian (78) retrieved ADRs associated with breast cancer drugs by mining user reviews from askapatient.com, drugs.com, and drugratingz.com and produced the ADRTrace medical term synonym set.

Extraction and prediction of ADRs from social media are challenging, especially because of inconsistency in the language used to report ADRs by different social groups, settings, and medical conditions (71). Indeed, various types of evidence exist in scientific publications (e.g., *in vitro*, *in vivo*, clinical) and social media (e.g., short sentences on Twitter, long comments on Instagram) to report ADRs and DDIs. To deal with this problem, data scientists in the field use both manual and automatic methods. The former involve manual curation by experts for each context, leading to the development of context-specific lexica and corpora, such as scientific literature reporting pharmacokinetics (83, 84) or pharmacogenetics (85, 86) studies, tweets mentioning medication intake (87), or Instagram user timelines annotated with standardized drug names and symptoms (3). There is also a corpus for comparative pharmacovigilance comprising 1,000 tweets and 1,000 PubMed sentences, with entities such as drugs, diseases, and symptoms (88). Such corpora are very useful for training automatic methods to identify pharmacological relevance in both social media and scientific literature.

Automatic methods to deal with language inconsistency include automatic topic modeling and word embedding techniques that cluster similar terms according to their co-occurrence patterns with other terms (89), typically implemented with spectral methods such as singular value decomposition (SVD) (90). More recently, word embeddings using neural networks, such as Word2vec (91), have shown much promise in obtaining high-quality word similarity spaces for biomedical text (92) and drug lexicons for social media analysis (93). Interestingly, SVD provides a linear approximation of, and insight into, what neural networks do in each layer (94) and a fast method to train them (95).

An example of automatic methods to deal with language inconsistency in pharmacovigilance for social media is ADRMine (96). It uses conditional random fields—a supervised sequence labeling classifier—to extract ADR mentions from tweets. The performance of the system is greatly enhanced by preprocessing posts (from Twitter and DailyStrength) for term similarity features using unsupervised word embeddings obtained via deep learning. Similarly, Word2vec embedding has been shown to increase the performance of automatic estimation of ADR rates of ten popular psychiatric drugs from Twitter, Reddit, and LiveJournal data (97), in comparison to the rates of ADRs for these drugs in the SIDER database (98). Interestingly, the lexicon derived by Word2vec leverages variants of ADR terms to deal with language inconsistency. A drawback of using deep learning methods, however, is the need for large training corpora, which limits the applicability to very commonly prescribed and discussed drugs and ADRs.

Most work using social media data for pharmacovigilance has focused on detecting signals for single drugs and their ADRs, although a few groups have studied DDI phenomena as well. Yang & Yang analyzed data from the patient discussion boards MedHelp, PatientsLikeMe, and DailyStrength and focused on 13 drugs and three DDI pairs (28). The study used association mining and DrugBank as a validation database (gold standard) with good results. From these datasets, the same group later built heterogeneous networks, where nodes represented such entities as “users,” “drugs,” or “ADR,” while edges signified “cause” or “treatment.” They went on to show that network motifs were effective in predicting DDIs for an expanded set of 23 drugs, using logistic regression as a link prediction classifier (99).

Soon after, Correia et al. (3) were the first to study DDI phenomena from all available posts on the popular social media site Instagram. The group focused on seven drugs known to treat depression, collected a large dataset of more than 5 million posts, and analyzed a population of almost

7 thousand users. Their study demonstrated the ability to analyze large cohorts of interest on popular social media sites. They also used a (heterogeneous) network science approach to produce a network of more than 600 drug, symptom, and natural product (NP) entities to monitor—via the web tool SyMPToM (Social Media Public healTh Monitoring; <https://symptom.sice.indiana.edu/>)—individuals and groups of patients, ADRs, DDIs, and conditions of interest. The top predicted links were validated against Drugbank and showed that the network approach allows for the identification and characterization of specific subcohorts (e.g., psoriasis patients and eating disorder groups) of relevance in the study of depression. Later on, the group expanded their work to include other epilepsy and opioid drugs, as well as analysis of Twitter data (65).

Recently, due to the opioid epidemic afflicting the United States, there has been an increased interest in using social media data to understand drug abuse (100). Several studies have analyzed licit (101) (chiefly alcohol), illicit (102) (e.g., cocaine and marijuana), and controlled substances (31, 103) (e.g., opioids) in diverse social media sites. Results are encouraging. For instance, analysis of Twitter data showed that geographical activity of posts mentioning prescription opioid misuse was strongly correlated with official government estimates (104), and deep learning methods could be used to predict opiate relapse using Reddit data (105). However, an older study that considered both questionnaires and Facebook data on five behavioral categories—including smoking, drinking, and illicit drug use—reported no significant correlation between respondents' Facebook profiles and illicit drug use (106). Analysis of such data on Internet health forums has also shown promise. Since these forums are often anonymous, open discussion about drug abuse may be more forthcoming. One study about the drug buprenorphine, a semisynthetic opioid effective in the treatment of opioid dependence, uncovered qualitative observations of public health interest such as increased discussion over time, perspectives on its utility, and reports of concomitant use with illicit drugs, which poses a significant health risk (107).

Social media data could also be useful in the study of the use, potential interactions, and effects of NPs and alternative medicines, including cannabis. Sales of NPs have nearly tripled over the past 20 years since passage of the Dietary Supplement Health and Education Act (108). Based on the general perception that “natural” means safe, the lay public often turn to NPs without discussing them with their healthcare practitioners (109). Consequently, patients frequently take NPs in conjunction with conventional medications, potentially triggering NP–drug interactions. The pharmacology of such products constitutes an array of ADRs and DDIs that have been very poorly explored by biomedical research (108). This is, thus, an arena where social media mining could provide important novel discoveries, early warnings, and insights, particularly with the possibility of studying large cohorts longitudinally (35). A preliminary study with NP and drug name dictionaries showed that it is possible to study their concomitant use longitudinally on Instagram and characterize associated symptoms with heterogeneous knowledge networks (3, 65).

3. CHARACTERIZING INDIVIDUAL AND COLLECTIVE PSYCHOLOGICAL WELL-BEING

The psychological and social well-being of individuals and populations is important, complex, and profoundly involved in shaping overall health-related phenomena. Scalable methodologies to gauge the changing mood of large populations from social media data—using NLP, sentiment analysis, ML, spectral methods, etc.—can help identify early warning indicators of lowered emotional resilience and potential health tipping points (e.g., the onset of mental disorders) for both epidemiological and precision health studies.

The brain is a complex system whose dynamics and function are shaped by the interactions of many components. Such systems can undergo critical transitions (CTs) (110), i.e., rapid and unexpected changes from one stable state to another, that are difficult to reverse. CTs provide

a powerful framework with which to understand and model mental health and its relation to the use of pharmaceuticals and other substances. For instance, recent clinical longitudinal studies have provided compelling evidence that psychological CTs, or tipping points, do occur in human psychological mood states. In particular, they are observed in the development of clinical depression and are preceded by measures of critical slowing down, such as increased autocorrelation and variance, as well as other common antecedents, that yield useful early-warning indicators of pending CTs (110, 111).

Social media data are a natural source of high-resolution, large-scale, longitudinal, introspective, and behavioral data to study, monitor, and even potentially intervene before CTs occur, avoiding significant hysteresis to return the system to a desirable state. For each individual social media user we can infer several important social and ethnographic factors from their online parameters (e.g., their location, language, and sex), as well as important risk information from their statements with respect to health-risk behavior and addiction and from their friendship and “follow” ties. In particular, it is now possible to widely track the evolving psychological mood state of social media users over extended periods of time along several relevant psychological dimensions (35, 112, 113). Indeed, social media indicators have been shown to predict the onset of depression (4, 33, 34, 114). Putting one’s own feelings into words on Twitter—also known as affect labeling—can sharply reverse negative emotions, demonstrating the attenuating effects of online affect labeling on emotional dynamics and its possible use as a mood regulation strategy (14) (see **Figure 2**). Measuring individual and collective sentiment from social media enables the design of actionable intervention strategies to alert individuals and communities to prevent the onset of mental health issues and health risk behavior [e.g., sexual activity (35)], especially in underserved or stigmatized populations (13) (see Section 4).

The term “sentiment analysis” refers to a set of computational techniques that are used to measure the opinions, sentiments, evaluations, appraisals, attitudes, and emotions that people express in natural language. This sentiment can be about entities such as products, services, organizations, individuals, issues, events, topics, and their respective attributes, but may also include self-referential elements (117, 118). Sentiment analysis is also broadly defined to include the computational treatment of opinion, mood, and subjectivity in text (119). The earliest studies of online sentiment relied on explicit user-defined features such as labels, ratings, reviews, etc. that were recorded as metadata to the text (118, 119). However, those features are not available for most online texts, including social media posts where health-related indicators need to be inferred from unstructured, unlabeled text. Indeed, the fundamental assumption of sentiment analysis is that individual- and population-level emotions are observable from unstructured written text.

Different methodological approaches have therefore been developed to extract sentiment indicators from text. Some methods use NLP and rely on the detection of word constructs (*n*-grams) in text to extract sentiment indicators with respect to an entity (120). Other techniques classify text into positive or negative mood classes using ML algorithms applied to annotated training sets, such as support vector machines (119) or naïve Bayes classifiers (112). Frequently, however, very good results are obtained with lexicon matching (119, 121, 122), a method that uses word lists (lexicons or dictionaries) of terms that are preannotated with sentiment values assigned by human subjects. Lexicons of sentiment-annotated terms are obtained via a variety of methods such as expert curation and consensus, population surveys, and automatic feature construction and selection in classification tasks (115, 118, 119, 121, 123–126). This approach is particularly useful when reliability over individual text segments is less important than scalability over large-scale datasets, as is the case for social media data.

Many sentiment lexicons focus on a single dimension of measured affect, such as negative or positive valence (roughly, happiness). Such lexicons include The General Inquirer (127), ANEW

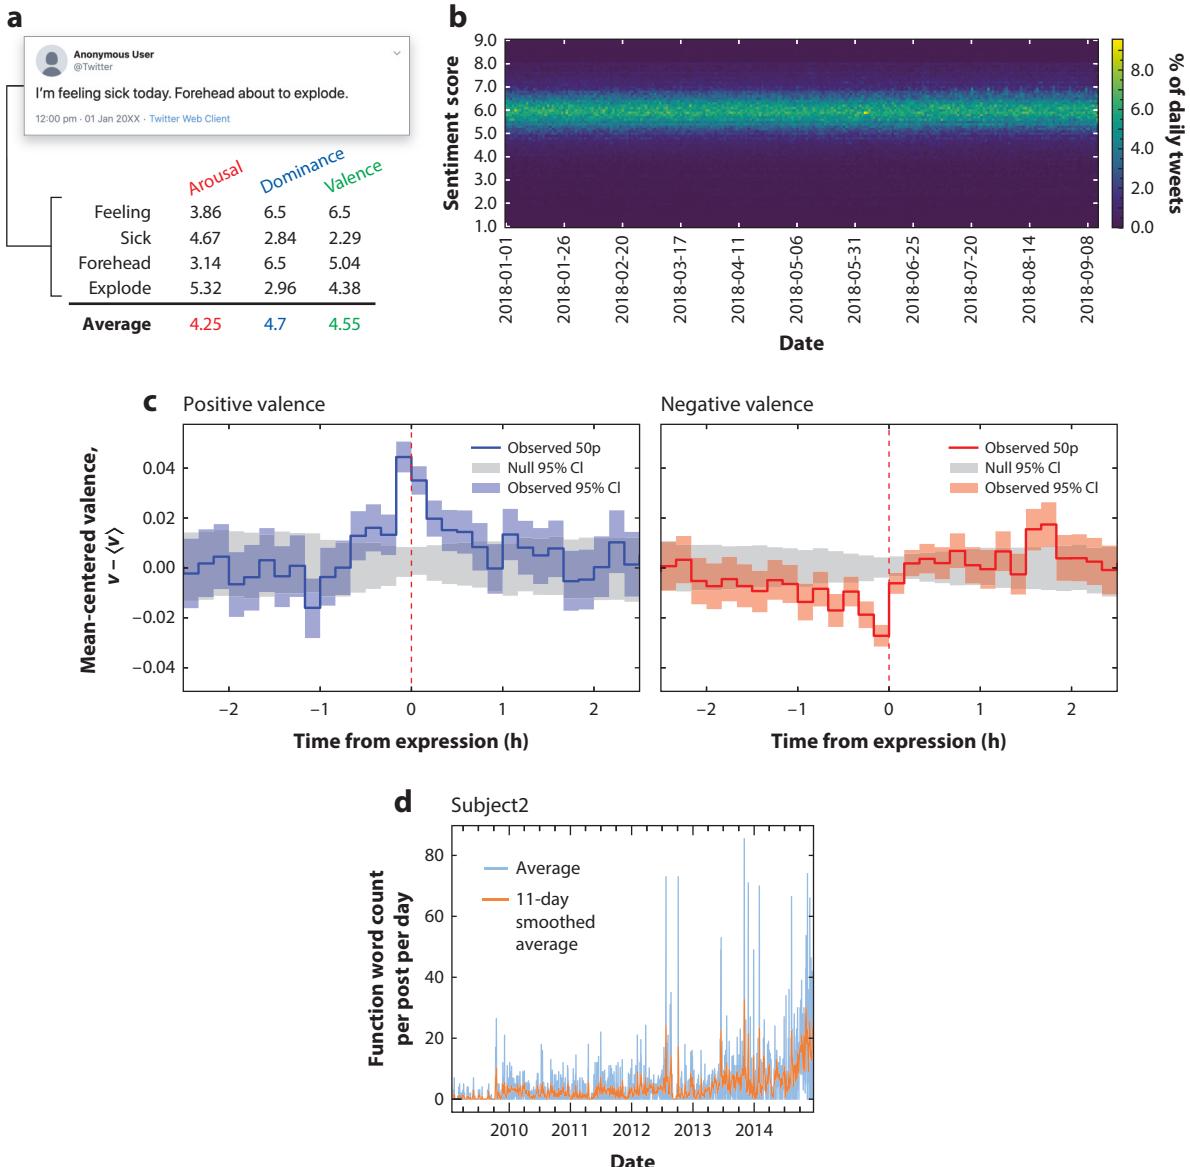


Figure 2

(a) An example tweet with its average ANEW (115) scores for arousal, dominance, and valence dimensions. Only words found in the ANEW dictionary were matched to their score. (b) A mood histogram time series showing the per-day distribution of ANEW valence scores for a cohort of Twitter users who self-reported being diagnosed with depression (116). (c) A mean-centered time series of ANEW valence scores for a cohort of Twitter users who stated that they were having a strong emotion on Twitter. Scores are shown for 1-min increments, smoothed by a 10-min rolling average, used to study (14) the effects of affect labeling on Twitter, i.e., the act of putting one's feeling into words, in this case by stating "I feel" in a tweet followed by a set of words that denote a strong emotion. Time $t = 0$ h (red dashed line) is the time at which the affect labeling tweet was posted for each person in the cohort. (d) Average LIWC (117) functional word count of the Facebook posts of a subject from a cohort of patients who died of SUDEP whose behavior on Facebook was studied after their death. This young patient, like several others in the cohort, showed an increase in functional words before SUDEP. Functional words are pronouns, prepositions, articles, conjunctions, auxiliary verbs, and a few other categories understood to indicate emotional states and other individual differences. Abbreviations: 50p, 50th percentile; ANEW, Affective Norms for English Words; CI, confidence interval; LIWC, Linguistic Inquiry and Word Count; SUDEP, sudden unexpected death in epilepsy.

(Affective Norms for English Words) (115, 121) along with its several extensions and translations (128), GPOMS (Google Profile of Mood States) (10, 126), LabMT (124), SentiWordNet (122), LIWC (Linguistic Inquiry and Word Count) (117), VADER (Valence-Aware Dictionary and Sentiment Reasoner) (129), and OpinionFinder (130). See the sidebar titled Sentiment Analysis Tools for more details about these tools and the categories or dimensions of sentiment that they aim to

SENTIMENT ANALYSIS TOOLS

The General Inquirer (127) is a tool to organize non-numerical data and to tag words in a text across various categories. The system started as a general-purpose tool with a dictionary of categories for the 3,000 most common English words and a few hundred words of interest to behavioral scientists. It has since grown to include the Harvard IV-4 and Lasswell content analysis dictionaries and has 198 categories.

ANEW (Affective Norms for English Words) includes ratings from 1 to 9 for 1,034 words along three mood dimensions: valence (unhappy to happy), arousal (calm to excited), and dominance (controlled to in control). These ratings are based on a nine-point, Likert-like scale and were collected from surveys given to psychology undergraduates (115). ANEW is used as a basis for several new dictionaries, including an extended version with nearly 14,000 words (128) and translations to several languages.

GPOMS (Google Profile of Mood States) is an extension of POMS (Profile of Mood States), a questionnaire of self-reported Likert scale questions measuring six underlying dimensions of mood: tension or anxiety; depression or dejection; anger or hostility; vigor or activity; fatigue or inertia; and confusion or bewilderment (133). GPOMS is used to translate this questionnaire into a dictionary suitable for sentiment analysis of large-scale social media data. Using word co-occurrences in Google's n -gram corpus, GPOMS extends the original 72 POMS terms to a dictionary of 964 terms that correspond to moods across six categories: calm, alert, sure, vital, kind, and happy (126). LabMT is a dictionary used by Amazon's Mechanical Turk to send out ANEW-like surveys that rank thousands of words on a nine-point scale from sad to happy and collects at least 50 ratings per word. Initially, LabMT comprised 10,222 English words found by merging the 5,000 most used words in each of four corpora: Google Books, Twitter, music lyrics, and *The New York Times* (124). This has since been extended to include 10 languages with about 10,000 words each collected across 24 corpora (123).

LIWC (Linguistic Inquiry and Word Count; pronounced "Luke") is a text analysis tool that has been actively supported and widely used since its first public release in 2001 (117, 134). LIWC was developed by several judges through a well-documented procedure that included external validity with psychological studies. The latest version, LIWC2015, has dictionaries containing nearly 6,400 words and produces outputs across about 90 categories including emotion and the use of pronouns, articles, cognitive processes, time focus, personal concerns, and informal language (117).

SentiWordNet (122) is a dictionary with word scores along three dimensions: positive, negative, and objective (neutral). It is built by automatically annotating sets of synonymous words (synsets) from WordNet (135), with additional steps using semisupervised classifiers and a random walk.

VADER (Valence-Aware Dictionary and Sentiment Reasoner) is a tool for measuring positive or negative text sentiment beyond simple lexicon matching (129). It searches for specific words in a sentence and modifies the associated dictionary-based sentiment scores based on simple rules, such as the presence of exclamations or negations. OpinionFinder (130) is a processing pipeline that tokenizes a document and then uses a series of classifiers trained on different corpora to find subjective statements and speech events. It identifies opinion sources and expressions of sentiment and classifies expressions as positive or negative.

In addition to these tools, other researchers suggest modifications of sentiment scores based on context by using compositional rules to modify sentiment scores from sentence parse trees (136, 137). Extensive reviews of sentiment analysis tools can be found in References 131 and 132.

capture. Many more sentiment analysis tools exist, with extensive reviews found in References 131 and 132.

These text analysis approaches, combined with large-scale social media data, have enabled the study of temporal patterns in the mood of populations at societal and global levels (35, 138). This includes studies of the changing features of language over time and geography (121, 125).

Because collective mood estimations are derived from tweets from large and diverse populations, the resulting distributions of sentiment values can contain distinct and informative components. However, many analyses of collective mood rely on average or median sentiment over time, which obscures this important information. Spectral methods such as SVD (90) have been effective in (a) removing the base sentiment contribution attributable to regular language use and (b) extracting sentiment components associated with specific phenomena of interest, e.g., moods correlated with increased interest in sex (35) or depression (116). These so-called eigen-moods are components that explain a significant proportion of the variation of sentiment in time series data instead of the average distribution of sentiment values that reflect prevailing language use. As such, they allow for more fine-grained assessments of individual- and population-level emotions associated with health behaviors of interest (116).

The different emotional dimensions of each sentiment analysis tool have been used for specific problems relevant to health and well-being. The authors of LIWC have demonstrated how its various indicators are useful to study the relation between language and a wide array of psychological problems (134). Underlying psychological states were shown to be revealed by various LIWC indicators, including the increased use of first-person singular pronouns to describe pain or trauma, the use of verb tenses to describe the immediacy of an experience, the use of first-person plural pronouns to denote higher social status, and prepositions and conjunctions used as proxies for thought complexity, among other examples, all of which enable the measurement of individual differences. For instance, textual features from the speech of student self-introductions measured by LIWC and analyzed with principal component analysis, were shown to be good predictors of overall academic performance. In particular, the use of commas, quotes, and negative affect were positively correlated with final performance, while the use of the present tense, first-person singular, and words from home, eating, and drinking categories were negatively correlated (139). LIWC has also been useful in classifying positive versus negative affect of dream reports (140) and completed versus noncompleted suicide attempts from suicide notes (141), as well as in measuring mood shifts and trends in large populations—e.g., feelings of sadness, anxiety, and anger—during extreme events like the September 11 World Trade Center attacks (142) or hurricanes (143). Diurnal and seasonal rhythms measured from Twitter data were found to be correlated with positive and negative sentiment as measured by LIWC. These measures thus revealed a change from positive to negative sentiment in the morning, as well as increased positive sentiment in days with more hours of daylight (125). Similarly, an analysis of more than 800 million Twitter posts for circadian mood variations further decomposed negative mood into anger, sadness, and fatigue, and found that fatigue follows an inverse pattern to the known circadian variation of concentrations of plasma cortisol—a hormone known to affect mood (144).

Many other sentiment analysis tools make use of lexicons. The ANEW lexicon (115, 128) consists of thousands of English words that have been rated by human subjects on three dimensions: valence, arousal, and dominance. This allows text sentiment to be analyzed along distinct emotional dimensions. ANEW has been used to show that the happiness of blogs steadily increased from 2005 to 2009, exhibiting a surprising increase in happiness with blogger age, but a decrease with distance from the Earth's equator (121). Twitter happiness has also been shown to follow a cycle that peaks on the weekend and bottoms out mid-week (121). The equally extensive LabMT lexicon has been used to demonstrate that sentiment measurements are robust to tuning

parameters and the removal of neutral terms (124). Indeed, a comparison of different sentiment analysis tools and their performance on several corpora, including LabMT, ANEW, and LIWC, found that these lexical tools tend to agree on positive and negative terms, but with notable differences in performance (131). Some lexicons are dedicated to specific application areas, e.g., subjective states (130), whereas others are geared toward general applicability. In general, lexical tools have been found to perform well only if the sentiment lexicon covers a large enough portion of the word frequency in a given text and its terms are scored on a continuous scale (132).

Social media data are also useful when sentiment analysis is applied to measure and address public health problems. Qualitative content analysis of sentiment (not using automatic sentiment analysis tools) on websites and discussion forums such as RateMDs.com has revealed a positivity bias in reviews of doctors (145) and that positive reviews are associated with surgeons who have a high volume of procedures (146). Similar qualitative content analysis applied to Twitter content has found mostly positive views of marijuana (147) with self-reports of personal use increasing when marijuana was legalized in two states (148). Most early sentiment studies of the relevance of social media for public health studies are based on qualitative, manual analysis, but there has been increased interest in large-scale, automatic studies. Using a custom sentiment analysis tool based on text classification (trained on annotated samples), Salathé & Khandelwal (36) studied dispositions toward flu vaccination on Twitter. They found that information flows more often between users who share the same sentiments and that most communities are dominated by either positive or negative sentiments toward a novel vaccine (homophily) (36). Unfortunately for public health campaigns, they also found that negative sentiment toward vaccines spreads more easily than positive sentiment in social networks (37).

Choudhury et al. (4) have shown that sentiment analysis tools like ANEW and LIWC are useful for analyzing the sentiment of tweets related to depression by building a large crowd-sourced corpus of tweets from individuals diagnosed with clinical depression (based on a standard psychometric instrument). They also introduced a social media depression index to characterize levels of depression in populations and demonstrated that its predictions were correlated with geographic, demographic, and seasonal patterns of depression reported by the Centers for Disease Control and Prevention (CDC). In addition to increased negative affect, onset of depression was also found to be correlated with a decrease in social activity, stronger clustering of social interactions, heightened relational and medicinal concerns, and greater expression of religious involvement (34).

Sentiment analysis of social media data was shown to help differentiate people on Twitter with posttraumatic stress disorder, depression, bipolar disorder, and seasonal affective disorder from control groups (149). This work also identified language and sentiment variables associated with each of the conditions. A similar study found that it is possible to identify linguistic and sentiment markers of schizophrenia on Twitter (150). Given that CTs in mental disease are likely associated with mood changes over time that can be captured by statistical parameters, like autocorrelation and variance (114), the multidimensional, large-scale data that can be extracted from social media, including sentiment, are likely to be of much use in the years to come.

ML methods such as deep learning have been used to accurately classify social media posts according to associated mental conditions (32). This has raised the possibility of characterizing a range of mental health conditions. The approach is still relatively new and most findings are preliminary (151), but it has stimulated an important discussion on the ethics of generating predictions about underlying conditions while also respecting health-related information privacy concerns.

4. OTHER PROMISING APPLICATIONS

Social media data have been used to study a wide range of other health-related problems and have yielded promising outcomes, especially when combined with other data sources. In disaster

and crisis informatics, when combined with physical sensor data, signals from social media have been useful in forecasting next-day smog-related health hazards (152). When used effectively by credible sources like emergency response teams, social media have also been used to mitigate community anxiety and the propagation of misinformation and rumors during and after environmental disasters (153).

In epidemiology, social media data have been useful in predicting disease outbreaks such as influenza (11, 23), cholera (24), and Zika (25). In the 2015–2016 Latin American outbreak of Zika, McGough et al. (25) used a dataset that combined Google searches and Twitter data to produce predictions of weekly suspected cases up to three weeks in advance of official publication. Such predictions have often been made based on correlations with the use of certain language, such as keywords or even emojis (154), or made indirectly through the use sentiment analysis tools (118). For instance, it has been shown that general Twitter sentiment about vaccines correlates with CDC estimates of vaccination rates by region (36) (see Section 3). Another study has shown that higher rates of tweets containing future-oriented language (e.g., “will,” “gonna”) are predictive of counties with lower HIV prevalence (26), demonstrating that social media may provide an inexpensive, real-time surveillance of disease outbreaks.

Social media research has also shown promise in efforts to combat stigma, offering a unique means by which to improve outcomes, benefiting healthcare providers and the public alike (155–157). Anti-stigma advocates and government organizations already have well-developed presences on Internet discussion boards and on social media for dozens of health conditions, including obesity/body issues (158) and HIV (159), which help raise awareness of health-related issues, including organ donations (160). Anti-stigma efforts around epilepsy, for example, include TalkAboutIt.org, a collaboration between actor Greg Grunberg and the Epilepsy Foundation (161, 162), and Out of the Shadows, a joint international project among the World Health Organization, the International League Against Epilepsy, and the International Bureau for Epilepsy. Efforts generally center around education about the disease, increasing awareness of epilepsy as a treatable brain disorder, and raising public acceptability of epilepsy. While these efforts do not utilize data science *per se*, engagement with social media platforms allows for data collection for future analysis of stigma in health.

There are not much data on the efficacy and long-term success of such anti-stigma efforts for epilepsy (163) or mental health disorders (164), although what exists has provided important insights. In their review, Patel et al. (165) have shown the benefits of social media anti-stigma efforts, with 48% of studies indicating positive results, 45% reporting results that were undefined or neutral, and 7% reporting potentially harmful effects. Researchers accessing Twitter have found 41% of epilepsy-related tweets to be derogatory (166). An analysis of the top ten epilepsy-related videos on YouTube has revealed that real-life or lived-experience videos garner the most hits, comments, and empathetic scores but provide little information. Videos with important health information, in contrast, have received only neutral or negative empathy scores (155). As contributing factors, concerns about privacy and the reactions of others limit respondents’ willingness to access and engage with content on a website (167). In one of the only network-based studies conducted so far, Silenzio et al. (168) found by mapping social network interactions that some individuals on social media may be more important to the spread of anti-stigma interventions. Additional research in this area is clearly needed.

The study of health-related issues around human sexuality can also be improved by analysis of web search, social media discourse, and health forum data, especially on those platforms that provide anonymity such as Reddit (45). For instance, web search and Twitter data have been instrumental in clarifying competing hypotheses about the cyclic sexual and reproductive behavior of human populations. Analysis of global data has suggested that rather than an evolutionary

adaptation to the annual solar cycle, observed birth cycles are likely a cultural phenomenon, since characteristic emotions around major cultural and religious celebrations (measured by sentiment analysis) correlate with interest in sex [measured by Google search data and birth records (35)].

On the medical side of human reproduction issues, pregnant women have frequently turned to the Internet and social media for reassurance on the normalcy of their pregnancies and to gather information on normal pregnancy symptoms, pregnancy complications, birth, and labor (169). For first-time mothers in particular, social media platforms have appeared to be the preferred mechanisms for obtaining important information during the antepartum and postpartum periods (170). Posting status updates and photos on social media appears to have provided pregnant women with a sense of connection with their peers, as well as with their own unborn babies (169, 171). Considering, in addition, the numbers of legal and illegal drug users on social media, as described above, social media platforms appear to be untapped sources of large-scale data on underreported population-level risk for neonatal and related conditions, such as neonatal abstinence syndrome. Social media signals may be effective resources to model the pharmacological, phenotypical, and psychosocial markers associated with drug use during pregnancy, and may lead to better early-problem warnings and prevention strategies.

Other measures that are known to correlate with health outcomes have also been investigated. For instance, social media deviations in diurnal rhythms, mobility patterns, and communication styles across regions have been included in a model that produces an accurate reconstruction of regional unemployment incidence (172). Additionally, the potential to use social media to predict severe health outcomes in epilepsy is preliminary but promising. Sudden unexpected death in epilepsy (SUDEP), for example, remains a leading cause of death in people with epilepsy. A small study of the Facebook timelines of people who died in this way was conducted to identify potential behavioral signs preceding SUDEP and has suggested that prior to dying a majority of the subjects wrote more text than they had previously on social media (see **Figure 2**).

5. LIMITATIONS

Social media can yield useful healthcare information, but there are inherent limitations to their use for biomedical applications. On the positive side, because analysis takes place after the data are recorded, social media analysis in general avoids experimenter and social conformity bias. Social media data constitute a type of real-world data (173) that allow for very large-scale population samples that surpass those of traditional social science and clinical trial approaches by several orders of magnitude. Indeed, Twitter offers strong opportunities for academic research given its public nature, real-time communication, and user population that approaches significant pluralities of the world's population. This is also the case for other social media platforms such as Facebook, Instagram, and Reddit. However, social media data frequently lack demographic indicators and ground truth, possibly resulting in biased or poorly representative samples—particularly when compared to the precisely defined inclusion and exclusion criteria of randomized controlled trials (RCTs) (173). In this section we provide a short overview of the literature related to the challenges of deriving valid and reliable indicators of human behavior from social media data and how these challenges can be mitigated.

In spite of scale, social media data generally entail self-selected samples, since subjects are free to choose when to participate and what content to submit. This bias is compounded by a mix of access restrictions imposed by social media platforms (174). As a result, researchers are prone to use so-called convenience samples, i.e., social media datasets that are, due to standardization efforts, more widespread, accessible, and convenient to use, although potentially not representative of the

wider population. Combined, these biases may lead to samples that do not validly or completely represent human behavior and diversity.

Social media content may also be subject to lexical bias (123) that could cause sentiment data to overrepresent positive sentiment. In addition, platform-specific factors may alter user behavior (174, 175) and lead to bias in subsequent data analysis. Indeed, users may be encouraged to engage in profile and reputation management by establishing different online personas to highlight their individuality and qualities that are perceived as desirable (176).

Privacy issues and algorithmic bias may also lead to mischaracterization of human factors. The behavior of most social media users is profoundly shaped by interface designs and, increasingly, algorithmic factors, e.g., the use of ML services for recommendations of social relations and relevant content. Nonhuman participants such as bots are, furthermore, widespread in some social media sites, e.g., Twitter. Moreover, exogenous events such as polarized elections may trigger individual and global sentiment changes, discourse polarization, and temporary deviations from baseline social linkage dynamics. The particular social-economic-political context in which social media data are recorded therefore plays an important role in analysis. Given these potential population biases, mining social media for healthcare information relevant to the broader human population requires a careful consideration of the multilevel complexity of human health (55), in which social and behavioral contexts play a critical role (177).

Perhaps one of the most important issues with social media mining is the difficulty of establishing sample validity and precise inclusion and exclusion criteria. Primarily, two sources of bias impact harvested social media data: sampling bias and algorithm bias. Sampling bias means that researchers cannot treat sampled social media data (e.g., a sample of 1% of tweets) as a truly representative and random sampling of the human population. This affects efforts to build valid cohorts and make generalizations from analyses (174). Samples cannot be balanced because of the lack of ground truth with respect to user demographics. Furthermore, the demographics of social media sites can vary broadly. In a survey of social media usage among American adults, 43% of women said they have used Instagram at least once, while for men this number was only 31%. Similarly, Hispanics appear underrepresented on LinkedIn—only 16% said they have used the platform as compared to 24% of Whites and 28% of Blacks. At the same time, Hispanics appear to be the largest demographic on WhatsApp—42% as compared to 24% of Whites and 13% of Blacks (178).

Sampling bias can be accentuated when subcohorts of social media users are used to draw geographical inferences, e.g., when particular key terms are used to infer location. Such subsamples may vary considerably in the degree to which they represent an unbiased sample. Future research using social media data must benefit from the large-scale nature of this real-world data, while specifying more precise inclusion and exclusion criteria, as used in RCTs, to avoid sample biases (173). Getting to that point requires the ability to stratify social media user cohorts using more fine-tuned ML, as well as via greater collaboration with and openness from social media platform providers. It is encouraging that Twitter data have been shown to match census and mobile phone data in geographical grids down to a square-kilometer resolution (179). Indeed, ML methods can be used on Twitter data to automatically track the incidence of health states in a population (180). Moreover, user demographics such as age and gender can be estimated from user content with reasonable accuracy (14).

In addition to sample bias, it is important to be aware of algorithmic biases that result from interface design, policies, and incentives associated with social media platforms. Since company revenues are tied to targeted advertisement, social media algorithms are tailored for navigation retention and profile building. These algorithms are highly dynamic, proprietary, and secret, all of which have consequences for research reproducibility (175). Most researchers, like users, are

largely unaware of how platforms filter timelines and other user information (181). Therefore, greater openness on the part of social media companies, perhaps encouraged or mandated by public policy, is needed to increase the utility of social media data for biomedicine.

In addition to sample and algorithmic biases, sentiment analysis can be manipulated by third parties through the injection of tweets (174), i.e., the deliberate insertion of tweets containing specific words known to affect sentiment tools, e.g., to boost general sentiment during a political debate. These efforts can be difficult to detect and mitigate since they affect the sample *a priori*, before a researcher can apply efforts to unbias their sample and address sample validity. Indeed, the extraction of emotional and social indicators from social media is fraught with difficulty. Users may indirectly disclose mood states, sentiment, health behavior, and diet, but rarely do so explicitly (14). Social media users furthermore favor an idiomatic style and vernacular that are difficult to analyze with traditional NLP tools and supervised ML algorithms. Applications of the latter are hampered by the lack of vetted ground truth datasets and the highly dynamic nature of underlying emotional processes. Additional difficulties in analyzing social media discourse include subjective opinion, the use of sarcasm (particularly toward the effects of specific drugs), and the polarity of a sentiment-laden word or phrase in context (119, 129). For instance, social media users may use the term “prozac” in a variety of idiosyncratic ways, but not necessarily because they are actually administering the drug.

Users revealing sensitive personal information about others and information pertinent to their social relationships raises serious privacy concerns. Indeed, data from eight to nine social media relations of an individual are sufficient to predict that individual’s characteristics just as well as their own information (182). In other words, privacy concerns are not just a matter of what users reveal about themselves but also a matter of what their social relations (unwittingly) reveal about them. Some users are aware of this phenomenon, which lowers their motivation and willingness to participate in studies using social media data (167).

Another limitation, finally, is the danger of overfitting in subsequent analysis. Because of data availability and privacy issues, information on specific cohorts is derived from indicators that are in turn derived from the content they generate. This will favor certain content and cohorts, possibly leading to models that overfit the data and generalize poorly (22, 183).

6. CONCLUSION

The studies reviewed in Section 2 show that social media users discuss a wide variety of personal and medical issues on social media platforms, e.g., their medical conditions, prognoses, medications, treatments, and quality of life, including improvements and adverse effects they experience from pharmacological treatments (3, 27, 52, 184). This collective discourse in turn can be monitored for early warnings of potential ADRs and to identify and characterize underreported, population-level pathology associated with therapies and DDIs that are most relevant to patients (3, 15, 27, 29). The new data-enabled modes of pharmacovigilance that social media afford are likely to be particularly relevant for patient-centered management and prevention of chronic diseases (185), such as epilepsy (186) and inflammatory diseases (52), which continue to be the chief healthcare problem in the United States (187). The inclusion of signals from, and engagement with, social media in patient-centered personal health library services that can store, recommend, and display individualized content to users is expected to significantly improve for chronic disease self-management (185), which is known to significantly lower disease severity and the number of unhealthy days and improve quality of life (188).

It is clear that disease prevention is increasingly becoming a matter of mitigating disease risk factors caused by an individual’s lifestyle and decision making, which are subject to a range of

cognitive, emotional, and social factors that have until now been difficult to assess with sufficient accuracy and scale. An understanding of the emotional and social factors that contribute to the emergence of public health issues is crucial for efficient mitigation strategies. As we describe in Section 3, there is already a substantial body of literature characterizing psychological well-being, especially by measuring individual and collective sentiment and other social interactions online. These methods have been particularly effective when used in combination with other sources of health and human behavior data, from physical sensors, mobility patterns, EHRs, and more precise physiological data.

The methodologies we cover also fall in the area of computational social science, which is presently focused on establishing the methodological framework to monitor societal phenomena from large-scale social media data—the abovementioned social macroscopes. For this methodology to be relevant in the prevention of disease and improvement of public health, researchers need to move from descriptive inductive modes of analysis to explanatory models with predictions and testable hypotheses. In particular, researchers need to establish social media not just as a tool for observation, but also as the foundation for explanatory models of the generative factors in health behavior and outcomes, of the type that computational and complexity sciences are already producing, e.g., in molecular and organismal biology (189).

There is reason to be optimistic about our ability to reach such predictive explanatory models since we know from psychological research that emotions play a significant role in human decision making (190). Behavioral finance in particular, for example, has provided evidence that financial decisions are significantly driven by emotion and mood (191). Hence it is reasonable to assume that online mood and sentiment, as well as all social media analysis we review, may be used to predict health behaviors and can therefore be used to predict individual and societal health outcomes.

The literature reviewed also points to a newfound ability to use social media data for improved well-being of small, specific cohorts and even individuals using precise characterizations and interventions. These may include pharmacological warnings, patient-centered management of chronic disease, and mental disorder assistance. For instance, donated timelines from individuals at risk of suicide can help ML models recognize early warning symptoms of depression and suicide (192).

Despite the proven importance to the specific goal of improving human health, social media data have been increasingly difficult to collect. Only a few social media data sources remain open for scientists. Many previously accessible sites are now almost entirely sealed from researchers, which is surprising given that the data are generated by and for its users, not the platforms, which mostly serve a mediating function. These limitations explain why most of the work reported has focused on Twitter, which remains open for data analysis. Nonetheless, other social networks have been shown to be useful for biomedicine, including Facebook (12), Instagram (3), Reddit (45–47), and even YouTube (48, 49).

It is possible that government policies may be leveraged to ensure accessibility to these important data sources, which could be considered a public good to be regulated much like publicly funded scientific publication data (193). This would help improve the sample and algorithmic limitations discussed in Section 5 and allow these large-scale, real-world data to better identify health factors that more expensive clinical trials cannot due to their smaller scale and cost. We intend for our review to contribute to establishing the importance of social media data for biomedical research and demonstrate the need to make such data more accessible in general to scientific research.

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

The authors would like to thank Deborah Rocha for editing this review and Marijn ten Thij for providing plots for figures. R.B.C. was funded by the CAPES Foundation (grant 18668127) and Fundação para a Ciência e a Tecnologia (grant PTDC-MEC-AND-30221-2017). J.B. thanks the Economic Development Agency (grant ED17HDQ3120040), the National Science Foundation (SMA/SME grant 1636636), the SparcS Center of Wageningen University, and the ISI Foundation for their support. L.M.R. was funded by the National Library of Medicine Program of the NIH (National Institutes of Health) (grants 1R01LM012832-01 and 01LM011945-01), by the NSF Research Traineeship Program (grant 1735095, “Interdisciplinary Training in Complex Networks and Systems”), and by Fundação Luso-Americana para o Desenvolvimento and the NSF (grant 276/2016). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

LITERATURE CITED

1. Correia RB, de Araújo Kohler LP, Mattos MM, Rocha LM. 2019. City-wide electronic health records reveal gender and age biases in administration of known drug–drug interactions. *NPJ Digit. Med.* 2:74
2. Christakis NA, Fowler JH. 2010. Social network sensors for early detection of contagious outbreaks. *PLOS ONE* 5:e12948
3. Correia RB, Li L, Rocha LM. 2016. Monitoring potential drug interactions and reactions via network analysis of Instagram user timelines. *Pac. Symp. Biocomput.* 21:492–503
4. Choudhury MD, Counts S, Horvitz E. 2013. Social media as a measurement tool of depression in populations. In *Proceedings of the 5th Annual ACM Web Science Conference*, pp. 47–56. New York: Assoc. Comput. Mach.
5. Lazer D, Pentland A, Adamic L, Aral S, Barabási AL, et al. 2009. Computational social science. *Science* 323:721–23
6. Salathé M, Bengtsson L, Bodnar TJ, Brewer DD, Brownstein JS, et al. 2012. Digital epidemiology. *PLOS Comput. Biol.* 8:e1002616
7. Varol O, Ferrara E, Ogan CL, Menczer F, Flammini A. 2014. Evolution of online user behavior during a social upheaval. In *Proceedings of the 2014 ACM Conference on Science*, pp. 81–90. New York: Assoc. Comput. Mach.
8. Kallus N. 2014. Predicting crowd behavior with big public data. In *WWW '14 Companion: Proceedings of the 23rd International Conference on the World Wide Web*, pp. 625–30. New York: Assoc. Comput. Mach.
9. Shao C, Hui PM, Wang L, Jiang X, Flammini A, et al. 2018. Anatomy of an online misinformation network. *PLOS ONE* 13:e0196087
10. Bollen J, Mao H, Zeng X. 2011. Twitter mood predicts the stock market. *J. Comput. Sci.* 2:1–8
11. Kautz H. 2013. *Data mining social media for public health applications*. Paper presented at the 23rd International Joint Conference on Artificial Intelligence (IJCAI 2013), Beijing, China, Aug. 5–9
12. Bakshy E, Messing S, Adamic LA. 2015. Exposure to ideologically diverse news and opinion on Facebook. *Science* 348:1130–32
13. Pescosolido BA, Martin JK. 2015. The stigma complex. *Annu. Rev. Sociol.* 41:87–116
14. Fan R, Varol O, Varamesh A, Barron A, van de Leemput IA, et al. 2018. The minute-scale dynamics of online emotions reveal the effects of affect labeling. *Nat. Hum. Behav.* 3:92–100
15. Paul MJ, Sarker A, Brownstein JS, Nikfarjam A, Scotch M, et al. 2016. Social media mining for public health monitoring and surveillance. *Pac. Symp. Biocomput.* 21:468–79
16. Hawn C. 2009. Take two aspirin and tweet me in the morning: how Twitter, Facebook, and other social media are reshaping health care. *Health Aff.* 28:361–68
17. Seltzer E, Jean N, Kramer-Golinkoff E, Asch D, Merchant R. 2015. The content of social media’s shared images about Ebola: a retrospective study. *Public Health* 129:1273–77
18. Sullivan R, Sarker A, O’Connor K, Goodin A, Karlsrud M, Gonzalez G. 2016. Finding potentially unsafe nutritional supplements from user reviews with topic modeling. *Pac. Symp. Biocomput.* 21:528–39

19. Hobbs WR, Burke M, Christakis NA, Fowler JH. 2016. Online social integration is associated with reduced mortality risk. *PNAS* 113:12980–84
20. Chan EH, Sahai V, Conrad C, Brownstein JS. 2011. Using web search query data to monitor dengue epidemics: a new model for neglected tropical disease surveillance. *PLOS Negl. Trop. Dis.* 5:e1206
21. Goel S, Hofman JM, Lahaie S, Pennock DM, Watts DJ. 2010. Predicting consumer behavior with web search. *PNAS* 107:17486–90
22. Lazer D, Kennedy R, King G, Vespignani A. 2014. The parable of Google Flu: traps in big data analysis. *Science* 343:1203–5
23. Signorini A, Segre AM, Polgreen PM. 2011. The use of Twitter to track levels of disease activity and public concern in the U.S. during the influenza a H1N1 pandemic. *PLOS ONE* 6:e19467
24. Chunara R, Andrews JR, Brownstein JS. 2012. Social and news media enable estimation of epidemiological patterns early in the 2010 Haitian cholera outbreak. *Am. J. Trop. Med. Hyg.* 86:39–45
25. McGough SF, Brownstein JS, Hawkins JB, Santillana M. 2017. Forecasting Zika incidence in the 2016 Latin America outbreak combining traditional disease surveillance with search, social media, and news report data. *PLOS Negl. Trop. Dis.* 11:e0005295
26. Ireland ME, Schwartz HA, Chen Q, Ungar LH, Albarracín D. 2015. Future-oriented tweets predict lower county-level HIV prevalence in the United States. *Health Psychol.* 34:1252–60
27. Hamed AA, Wu X, Erickson R, Fandy T. 2015. Twitter K-H networks in action: advancing biomedical literature for drug search. *J. Biomed. Inform.* 56:157–68
28. Yang H, Yang CC. 2013. Harnessing social media for drug-drug interactions detection. In *2013 IEEE International Conference on Healthcare Informatics*, pp. 22–29. New York: IEEE
29. Sarker A, Ginn R, Nikfarjam A, O'Connor K, Smith K, et al. 2015. Utilizing social media data for pharmacovigilance: a review. *J. Biomed. Inform.* 54:202–12
30. Pain J, Levacher J, Quinquenel A, Belz A. 2016. Analysis of Twitter data for postmarketing surveillance in pharmacovigilance. In *Proceedings of the 2nd Workshop on Noisy User-generated Text*, pp. 94–101. Osaka, Jpn.: COLING 2016 Organ. Commit.
31. Sarker A, O'Connor K, Ginn R, Scotch M, Smith K, et al. 2016. Social media mining for toxicovigilance: automatic monitoring of prescription medication abuse from Twitter. *Drug Saf.* 39:231–40
32. Gkotsis G, Oellrich A, Velupillai S, Liakata M, Hubbard TJP, et al. 2017. Characterisation of mental health conditions in social media using informed deep learning. *Sci. Rep.* 7:45141
33. Eichstaedt JC, Smith RJ, Merchant RM, Ungar LH, Crutchley P, et al. 2018. Facebook language predicts depression in medical records. *PNAS* 115:11203–8
34. De Choudhury M, Gamon M, Counts S, Horvitz E. 2013. Predicting depression via social media. In *Proceedings of the Seventh International AAAI Conference on Weblogs and Social Media*, pp. 128–37. Menlo Park, CA: Assoc. Adv. Artif. Intell.
35. Wood IB, Varela PL, Bollen J, Rocha LM, Gonçalves-Sá J. 2017. Human sexual cycles are driven by culture and match collective moods. *Sci. Rep.* 7:17973
36. Salathé M, Khandelwal S. 2011. Assessing vaccination sentiments with online social media: implications for infectious disease dynamics and control. *PLOS Comput. Biol.* 7:e1002199
37. Salathé M, Vu DQ, Khandelwal S, Hunter DR. 2013. The dynamics of health behavior sentiments on a large online social network. *EPJ Data Sci.* 2:4
38. Statista. 2019. *Number of monthly active Twitter users worldwide from 1st quarter 2010 to 1st quarter 2019 (in millions)*. Social Media Stat., accessed Nov. 29. <https://www.statista.com/statistics/282087/number-of-monthly-active-twitter-users/>
39. Statista. 2019. *Number of monthly active Instagram users from January 2013 to June 2018 (in millions)*. Social Media Stat., accessed Nov. 29. <https://www.statista.com/statistics/253577/number-of-monthly-active-instagram-users/>
40. Statista. 2019. *Most popular social networks of teenagers in the United States from Fall 2012 to Spring 2019*. Social Media Stat., accessed Nov. 29. <https://www.statista.com/statistics/250172/social-network-usage-of-us-teens-and-young-adults/>
41. Statista. 2019. *Number of monthly active Facebook users worldwide as of 3rd quarter 2019 (in millions)*. Social Media Stat., accessed Nov. 29. <https://www.statista.com/statistics/264810/number-of-monthly-active-facebook-users-worldwide/>

42. Statista. 2019. *Facebook—statistics & facts*. Social Media Stat., accessed Nov. 29. <https://www.statista.com/topics/751/facebook/>
43. NIH (Natl. Inst. Health). 2014. *Early stage development of technologies in biomedical computing, informatics, and big data science (R01)*. Funding Oppor. Announc., Natl. Inst. Health, Bethesda, MD. <https://grants.nih.gov/grants/guide/pa-files/PA-14-155.html> 29
44. NIH (Natl. Inst. Health). 2014. *Extended development, hardening and dissemination of technologies in biomedical computing, informatics and big data science (R01)*. Funding Oppor. Announc., Natl. Inst. Health, Bethesda, MD. <https://grants.nih.gov/grants/guide/pa-files/PA-14-156.html>
45. Choudhury MD, De S. 2014. Mental health discourse on reddit: self-disclosure, social support, and anonymity. In *Proceedings of the Eighth International AAAI Conference on Weblogs and Social Media*, pp. 71–80. Menlo Park, CA: Assoc. Adv. Artif. Intell.
46. Park A, Conway M. 2017. Tracking health related discussions on Reddit for public health applications. *AMIA Annu. Symp. Proc.* 2017:1362–71
47. Zomick J, Levitan SI, Serper M. 2019. Linguistic analysis of schizophrenia in Reddit posts. In *Proceedings of the Sixth Workshop on Computational Linguistics and Clinical Psychology*, pp. 74–83. Minneapolis, MN: Assoc. Comput. Linguist.
48. Fernandez-Luque L, Elahi N, Grajales FJ 3rd. 2009. An analysis of personal medical information disclosed in YouTube videos created by patients with multiple sclerosis. In *Medical Informatics in a United and Healthy Europe*, ed. K-P Adlassnig, B Blobel, J Mantas, I Masic, pp. 292–96. Amsterdam: IOS
49. Syed-Abdul S, Fernandez-Luque L, Jian WS, Li YC, Crain S, et al. 2013. Misleading health-related information promoted through video-based social media: anorexia on YouTube. *J. Med. Internet Res.* 15:e30
50. Pescosolido BA, Olafsdottir S, Sporns O, Perry BL, Meslin EM, et al. 2016. The social symbiome framework: linking genes-to-global cultures in public health using network science. In *Handbook of Applied System Science*, pp. 25–48, ed. ZP Neal. New York: Routledge
51. Fernández-Luque L, Bau T. 2015. Health and social media: perfect storm of information. *Healthc. Inform. Res.* 21:67–73
52. Patel R, Belousov M, Jani M, Dasgupta N, Winakor C, et al. 2018. Frequent discussion of insomnia and weight gain with glucocorticoid therapy: an analysis of Twitter posts. *NPJ Dig. Med.* 1:20177
53. Cooper V, Metcalf L, Versnel J, Upton J, Walker S, Horne R. 2015. Patient-reported side effects, concerns and adherence to corticosteroid treatment for asthma, and comparison with physician estimates of side-effect prevalence: a UK-wide, cross-sectional study. *NPJ Prim. Care Respir. Med.* 25:15026
54. Börner K. 2011. Plug-and-play macroscopes. *Commun. ACM* 54:60–69
55. Alber M, Tepole AB, Cannon WR, De S, Dura-Bernal S, et al. 2019. Integrating machine learning and multiscale modeling—perspectives, challenges, and opportunities in the biological, biomedical, and behavioral sciences. *NPJ Dig. Med.* 2:115
56. Sultana J, Cutroneo P, Trifirò G. 2013. Clinical and economic burden of adverse drug reactions. *J. Pharmacol. Pharmacother.* 4:S73–77
57. Becker ML, Kallewaard M, Caspers PW, Visser LE, Leufkens HG, Stricker BH. 2007. Hospitalisations and emergency department visits due to drug-drug interactions: a literature review. *Pharmacoepidemiol. Drug Saf.* 16:641–51
58. FDA (US Food Drug Admin.). 2019. *Questions and answers on FDA's Adverse Event Reporting System (FAERS)*. Fact Sheet, US Food Drug Admin., Silver Spring, MD. <https://www.fda.gov/drugs-surveillance/questions-and-answers-fdas-adverse-event-reporting-system-faers>
59. Alatawi YM, Hansen RA. 2017. Empirical estimation of under-reporting in the U.S. Food and Drug Administration Adverse Event Reporting System (FAERS). *Expert Opin. Drug Saf.* 16:761–67
60. Basch E. 2010. The missing voice of patients in drug-safety reporting. *N. Engl. J. Med.* 362:865–69
61. Rao AL, Hong ES. 2016. Understanding depression and suicide in college athletes: emerging concepts and future directions. *Br. J. Sports Med.* 50:136–37
62. Druckman JN, Rothschild JE. 2018. Playing with pain: social class and pain reporting among college student-athletes. *Sport J.* 21. <https://thesportjournal.org/article/playing-with-pain-social-class-and-pain-reporting-among-college-student-athletes/>

63. White R, Harpaz R, Shah N, DuMouchel W, Horvitz E. 2014. Toward enhanced pharmacovigilance using patient-generated data on the internet. *Clin. Pharmacol. Ther.* 96:239–46
64. Topaz M, Lai K, Dhopeshwarkar N, Seger DL, Sa'adon R, et al. 2016. Clinicians' reports in electronic health records versus patients' concerns in social media: a pilot study of adverse drug reactions of aspirin and atorvastatin. *Drug Saf.* 39:241–50
65. Correia RB. 2019. *Prediction of drug interaction and adverse reactions, with data from electronic health records, clinical reporting, scientific literature, and social media, using complexity science methods*. PhD Thesis, Luddy Sch. Inform. Comput. Eng., Indiana Univ., Bloomington, IN
66. Lardon J, Abdellaoui R, Bellet F, Asfari H, Souvignet J, et al. 2015. Adverse drug reaction identification and extraction in social media: a scoping review. *J. Med. Internet Res.* 17:e171
67. Ofoghi B, Mann M, Verspoor K. 2016. Towards early discovery of salient health threats: a social media emotion classification technique. *Pac. Symp. Biocomput.* 21:504–15
68. Weber I, Achananuparp P. 2016. Insights from machine-learned diet success prediction. *Pac. Symp. Biocomput.* 21:540–51
69. Aphinyanaphongs Y, Lulejian A, Brown DP, Bonneau R, Krebs P. 2016. Text classification for automatic detection of e-cigarette use and use for smoking cessation from Twitter: a feasibility pilot. *Pac. Symp. Biocomput.* 21:480–91
70. Sap M, Kern ML, Eichstaedt JC, Kapelner A, Agrawal M, et al. 2016. Predicting individual well-being through the language of social media. *Pac. Symp. Biocomput.* 21:516–27
71. Sarker A, Gonzalez-Hernandez G. 2017. Overview of the Second Social Media Mining for Health (SMM4H) shared tasks at AMIA 2017. In *Proceedings of the 2nd Social Media Mining for Health Research and Applications Workshop Co-Located with the American Medical Informatics Association Annual Symposium (AMIA 2017)*, ed. A Sarker, G Gonzalez, pp. 43–48. <http://ceur-ws.org/Vol-1996/paper8.pdf>
72. Weissenbacher D, Sarker A, Magge A, O'Connor ADK, Paul M, Gonzalez-Hernandez G. 2019. Overview of the Fourth Social Media Mining for Health (#SMM4H) shared task at ACL 2019. In *Proceedings of the 4th Social Media Mining for Health Applications (#SMM4H) Workshop & Shared Task*, pp. 21–30. Stroudsburg, PA: Assoc. Comput. Linguist.
73. Chee B, Karahalios KG, Schatz B. 2009. Social visualization of health messages. In *Proceedings of the 42nd Hawaii International Conference on System Sciences*. Los Alamitos, CA: IEEE Comput. Soc.
74. Leaman R, Wojtulewicz L, Sullivan R, Skariah A, Yang J, Gonzalez G. 2010. 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*, pp. 117–25. Stroudsburg, PA: Assoc. Comput. Linguist.
75. Nikfarjam A, Gonzalez GH. 2011. Pattern mining for extraction of mentions of adverse drug reactions from user comments. *AMIA Annu. Symp. Proc.* 2011:1019–26
76. Benton A, Ungar L, Hill S, Hennessy S, Mao J, et al. 2011. Identifying potential adverse effects using the web: a new approach to medical hypothesis generation. *J. Biomed. Inform.* 44:989–96
77. Sampathkumar H, Luo B, Chen XW. 2012. Mining adverse drug side-effects from online medical forums. In *2012 IEEE Second International Conference on Healthcare Informatics, Imaging and Systems Biology (HISB)*, Pap. 150. New York: IEEE
78. Yates A, Goharian N. 2013. ADRTrace: detecting expected and unexpected adverse drug reactions from user reviews on social media sites. In *Advances in Information Retrieval*, ed. P Serdyukov, P Braslavski, SO Kuznetsov, J Kamps, S Rüger, et al., pp. 816–19. Berlin: Springer-Verlag
79. Patki A, Sarker A, Pimpalkhute P, Nikfarjam A, Ginn R, et al. 2014. *Mining adverse drug reaction signals from social media: going beyond extraction*. Paper presented at BioLINK SIG 2014, Boston, MA, July 11–12
80. Chee B, Berlin R, Schatz B. 2009. Measuring population health using personal health messages. *AMIA Annu. Symp. Proc.* 2009:92–96
81. Yang CC, Yang H, Jiang L, Zhang M. 2012. Social media mining for drug safety signal detection. In *Proceedings of the 2012 International Workshop on Smart Health and Wellbeing*, pp. 33–40. New York: Assoc. Comput. Mach.
82. Yang M, Kiang M, Shang W. 2015. Filtering big data from social media—building an early warning system for adverse drug reactions. *J. Biomed. Inform.* 54:230–40

83. Wu HY, Karnik S, Subhadarshini A, Wang Z, Philips S, et al. 2013. An integrated pharmacokinetics ontology and corpus for text mining. *BMC Bioinform.* 14:35
84. Kolchinsky A, Lourenço A, Wu H-Y, Li L, Rocha LM. 2015. Extraction of pharmacokinetic evidence of drug–drug interactions from the literature. *PLOS ONE* 10:e0122199
85. Zhang P, Wu H, Chiang C, Wang L, Binkheder S, et al. 2018. Translational biomedical informatics and pharmacometrics approaches in the drug interactions research. *CPT Pharmacometr. Syst. Pharmacol.* 7:90–102
86. Wu HY, Shendre A, Zhang S, Zhang P, Wang L, et al. 2019. Translational knowledge discovery between drug interactions and pharmacogenetics. *Clin. Pharmacol. Ther.* 107:886–902
87. Klein A, Sarker A, Rouhizadeh M, O'Connor K, Gonzalez G. 2017. Detecting personal medication intake in Twitter: an annotated corpus and baseline classification system. In *Proceedings of the BioNLP 2017 Workshop*, pp. 136–42. Stroudsburg, PA: Assoc. Comput. Linguist.
88. Alvaro N, Miyao Y, Collier N. 2017. TwiMed: Twitter and Pubmed comparable corpus of drugs, diseases, symptoms, and their relations. *JMIR Public Health Surveill.* 3:e24
89. Blei DM, Lafferty JD. 2009. Topic models. In *Text Mining*, pp. 71–94. Boca Raton, FL: Chapman & Hall
90. Wall ME, Rechtsteiner A, Rocha LM. 2003. Singular value decomposition and principal component analysis. In *A Practical Approach to Microarray Data Analysis*, ed. DP Berrar, W Dubitzky, M Granzow, pp. 91–109. New York: Springer
91. Goldberg Y, Levy O. 2014. word2vec explained: deriving Mikolov et al.'s negative-sampling word-embedding method. arXiv:1402.3722 [cs.CL]
92. Jiang Z, Li L, Huang D, Jin L. 2015. Training word embeddings for deep learning in biomedical text mining tasks. In *2015 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, ed. J Huan, IM Schapranow, S Miyano, I Yoo, A Shehu, et al., pp. 625–28. New York: IEEE
93. Lavertu A, Altman RB. 2019. RedMed: extending drug lexicons for social media applications. *J. Biomed. Inform.* 99:103307
94. Bermeitinger B, Hrycej T, Handschuh S. 2019. Singular value decomposition and neural networks. In *Artificial Neural Networks and Machine Learning—ICANN 2019*, ed. IV Tetko, V Kürková, P Karpov, F Theis, pp. 153–64. Cham, Switz.: Springer
95. Cai C, Ke D, Xu Y, Su K. 2014. Fast learning of deep neural networks via singular value decomposition. In *PRICAI 2014: Trends in Artificial Intelligence*, ed. DN Pham, SB Park, pp. 820–26. Cham, Switz.: Springer
96. Nikfarjam A, Sarker A, O'Connor K, Ginn R, Gonzalez G. 2015. Pharmacovigilance from social media: mining adverse drug reaction mentions using sequence labeling with word embedding cluster features. *J. Am. Med. Inform. Assoc.* 22:671–81
97. Nguyen T, Larsen ME, O'Dea B, Phung D, Venkatesh S, Christensen H. 2017. Estimation of the prevalence of adverse drug reactions from social media. *Int. J. Med. Inform.* 102:130–37
98. Kuhn M, Letunic I, Jensen LJ, Bork P. 2016. The SIDER database of drugs and side effects. *Nucleic Acids Res.* 44:D1075–79
99. Yang H, Yang CC. 2015. Mining a weighted heterogeneous network extracted from healthcare-specific social media for identifying interactions between drugs. In *2015 IEEE International Conference on Data Mining Workshop (ICDMW)*, pp. 196–203. New York: IEEE
100. Kim SJ, Marsch LA, Hancock JT, Das AK. 2017. Scaling up research on drug abuse and addiction through social media big data. *J. Med. Internet Res.* 19:e353
101. West JH, Hall PC, Hanson CL, Prier K, Giraud-Carrier C, et al. 2012. Temporal variability of problem drinking on Twitter. *Open J. Prev. Med.* 2:43–48
102. Yakushev A, Mityagin S. 2014. Social networks mining for analysis and modeling drugs usage. *Procedia Comput. Sci.* 29:2462–71
103. Shutler L, Nelson LS, Portelli I, Blachford C, Perrone J. 2015. Drug use in the Twittersphere: qualitative contextual analysis of tweets about prescription drugs. *J. Addict. Dis.* 34:303–10
104. Chary M, Genes N, Giraud-Carrier C, Hanson C, Nelson LS, Manini AF. 2017. Epidemiology from tweets: estimating misuse of prescription opioids in the USA from social media. *J. Med. Toxicol.* 13:278–86
105. Yang Z, Nguyen L, Jin F. 2018. Predicting opioid relapse using social media data. arXiv:1811.12169 [cs.SI]

106. van Hoof JJ, Bekkers J, van Vuuren M. 2014. Son, you're smoking on Facebook! College students' disclosures on social networking sites as indicators of real-life risk behaviors. *Comput. Hum. Behav.* 34:249–57
107. Daniulaityte R, Carlson R, Brigham G, Cameron D, Sheth A. 2015. "Sub is a weird drug:" a Web-based study of lay attitudes about use of buprenorphine to self treat opioid withdrawal symptoms. *Am. J. Addict.* 24:403–9
108. Brantley SJ, Argikar AA, Lin YS, Nagar S, Paine MF. 2014. Herb–drug interactions: challenges and opportunities for improved predictions. *Drug Metab. Dispos.* 42:301–17
109. Blendon RJ, DesRoches CM, Benson JM, Brodie M, Altman DE. 2001. Americans' views on the use and regulation of dietary supplements. *Arch. Intern. Med.* 161:805–10
110. Scheffer M, Carpenter SR, Lenton TM, Bascompte J, Brock W, et al. 2012. Anticipating critical transitions. *Science* 338:344–48
111. van de Leemput IA, Wichers M, Cramer AOJ, Borsboom D, Tuerlinckx F, et al. 2014. Critical slowing down as early warning for the onset and termination of depression. *PNAS* 111:87–92
112. Pak A, Paroubek P. 2010. Twitter as a corpus for sentiment analysis and opinion mining. In *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10)*, pp. 1320–26. Paris: Eur. Lang. Resour. Assoc.
113. Zimbra D, Abbasi A, Zeng D, Chen H. 2018. The state-of-the-art in Twitter sentiment analysis: a review and benchmark evaluation. *ACM Trans. Manag. Inform. Syst.* 9:5
114. Reece AG, Reagan AJ, Lix KLM, Dodds PS, Danforth CM, Langer EJ. 2017. Forecasting the onset and course of mental illness with Twitter data. *Sci. Rep.* 7:13006
115. Bradley MM, Lang PJ. 1999. *Affective Norms for English Words (ANEW): instruction manual and affective ratings*. Instr. Man., Natl. Inst. Mental Health Cent. Stud. Emot. Atten., Univ. Florida, Gainesville, FL
116. ten Thij M, Bollen J, Rocha LM. 2019. Detecting eigenmoods in individual human emotion. In *Book of Abstracts of the 8th International Conference on Complex Networks and their Applications*, ed. H Cherifi, S Gaito, J Gonçalves-Sá, J Fernando Mendes, E Moro, et al., pp. 166–68. Lisbon, Port.: Int. Conf. Complex Netw. Appl.
117. Pennebaker JW, Boyd RL, Jordan K, Blackburn K. 2015. *The development and psychometric properties of LIWC2015*. Psychom. Man., Dep. Psych., Univ. Texas, Austin
118. Liu B. 2012. *Sentiment Analysis and Opinion Mining*. Williston, VT: Morgan & Claypool
119. Pang B, Lee L. 2008. Opinion mining and sentiment analysis. *Found. Trends Inform. Retr.* 2:1–135
120. Nasukawa T, Yi J. 2003. Sentiment analysis: capturing favorability using natural language processing. In *K-CAP '03: Proceedings of the 2nd International Conference on Knowledge Capture*, pp. 70–77. New York: Assoc. Comput. Mach.
121. Dodds PS, Danforth CM. 2010. Measuring the happiness of large-scale written expression: songs, blogs, and presidents. *J. Happiness Stud.* 11:441–56
122. Esuli A, Sebastiani F. 2006. SENTIWORDNET: a publicly available lexical resource for opinion mining. In *Proceedings of the Fifth International Conference on Language Resources and Evaluation (LREC'06)*, pp. 417–22. Paris: Eur. Lang. Resour. Assoc.
123. Dodds PS, Clark EM, Desu S, Frank MR, Reagan AJ, et al. 2015. Human language reveals a universal positivity bias. *PNAS* 112:2389–94
124. Dodds PS, Harris KD, Kloumann IM, Bliss CA, Danforth CM. 2011. Temporal patterns of happiness and information in a global social network: hedonometrics and Twitter. *PLOS ONE* 6:e26752
125. Golder SA, Macy MW. 2011. Diurnal and seasonal mood vary with work, sleep, and daylength across diverse cultures. *Science* 333:1878–81
126. Bollen J, Pepe A, Mao H. 2011. Modeling public mood and emotion: Twitter sentiment and socio-economic phenomena. In *Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media (ICWSM 2011)*, pp. 450–53. Menlo Park, CA: Assoc. Adv. Artif. Intell.
127. Stone PJ, Bales RF, Namenwirth JZ, Ogilvie DM. 1962. The General Inquirer: a computer system for content analysis and retrieval based on the sentence as a unit of information. *Behav. Sci.* 7:484–98
128. Warriner AB, Kuperman V, Brysbaert M. 2013. Norms of valence, arousal, and dominance for 13,915 English lemmas. *Behav. Res. Methods* 45:1191–207

129. Hutto C, Gilbert E. 2014. VADER: a parsimonious rule-based model for sentiment analysis of social media text. In *Proceedings of the Eighth International AAAI Conference on Weblogs and Social Media*, pp. 216–25. Menlo Park, CA: Assoc. Adv. Artif. Intell.
130. Wilson T, Hoffmann P, Somasundaran S, Kessler J, Wiebe J, et al. 2005. OpinionFinder: a system for subjectivity analysis. In *Proceedings of HLT/EMNLP 2005 on Interactive Demonstrations*, pp. 34–35. Stroudsburg, PA: Assoc. Comput. Linguist.
131. Ribeiro FN, Araújo M, Gonçalves P, Gonçalves MA, Benevenuto F. 2016. SentiBench: a benchmark comparison of state-of-the-practice sentiment analysis methods. *EPJ Data Sci.* 5:23
132. Reagan AJ, Danforth CM, Tivnan B, Williams JR, Dodds PS. 2017. Benchmarking sentiment analysis methods for large-scale texts: a case for using continuum-scored words and word shift graphs. *EPJ Data Sci.* 6:28
133. Douglas M, McNair ML, Droppleman LF. 1971. *POMS manual for the Profile of Mood States*. Instr. Man., Educ. Indust. Test. Serv., San Diego, CA
134. Tausczik YR, Pennebaker JW. 2010. The psychological meaning of words: LIWC and computerized text analysis methods. *J. Lang. Soc. Psychol.* 29:24–54
135. Miller GA. 1995. WordNet: a lexical database for English. *Commun. ACM* 38:39–41
136. Moilanen K, Pulman S. 2007. Sentiment composition. In *Proceedings of the Fourth International Conference on Recent Advances in Natural Language Processing (RANLP 2007)*, pp. 378–82. Stroudsburg, PA: Assoc. Comput. Linguist.
137. Saif H, He Y, Alani H. 2012. Semantic sentiment analysis of Twitter. In *Proceedings of the 11th International Semantic Web Conference ISWC*, pp. 508–24. Berlin: Springer
138. Hannak A, Anderson E, Barrett LF, Lehmann S, Mislove A, Riedewald M. 2012. Tweetin' in the rain: exploring societal-scale effects of weather on mood. In *Proceedings of the Sixth International AAAI Conference on Weblogs and Social Media*, pp. 479–82. Menlo Park, CA: Assoc. Adv. Artif. Intell.
139. Robinson RL, Navea R, Ickes W. 2013. Predicting final course performance from students' written self-introductions: a LIWC analysis. *J. Lang. Soc. Psychol.* 32:469–79
140. Nadeau D, Sabourin C, Koninek JD, Matwin S, Turney PD. 2006. *Automatic dream sentiment analysis*. Poster presented at Workshop on Computational Aesthetics at the Twenty-First National Conference on Artificial Intelligence, Boston, MA, July 16
141. Pestian JP, Matykievicz P, Linn-Gust M, South B, Uzuner O, et al. 2012. Sentiment analysis of suicide notes: a shared task. *Biomed. Inform. Insights* 5:3–16
142. Back MD, Küfner AC, Egloff B. 2011. “Automatic or the people?” Anger on September 11, 2001, and lessons learned for the analysis of large digital data sets. *Psychol. Sci.* 22:837–38
143. Kryvasheyeu Y, Chen H, Moro E, Hentenryck PV, Cebrian M. 2015. Performance of social network sensors during hurricane sandy. *PLOS ONE* 10:e0117288
144. Dzogang F, Lightman S, Cristianini N. 2017. Circadian mood variations in Twitter content. *Brain Neurosci. Adv.* <https://doi.org/10.1177/2398212817744501>
145. López A, Detz A, Ratanawongs N, Sarkar U. 2012. What patients say about their doctors online: a qualitative content analysis. *J. Gen. Intern. Med.* 27:685–92
146. Segal J, Sacopulos M, Sheets V, Thurston I, Brooks K, Puccia R. 2012. Online doctor reviews: Do they track surgeon volume, a proxy for quality of care? *J. Med. Internet Res.* 14:e50
147. Cavazos-Rehg PA, Krauss M, Fisher SL, Salyer P, Grucza RA, Bierut LJ. 2015. Twitter chatter about marijuana. *J. Adolesc. Health* 56:139–45
148. Thompson L, Rivara FP, Whitehill JM. 2015. Prevalence of marijuana-related traffic on Twitter, 2012–2013: a content analysis. *Cyberpsychol. Behav. Soc. Netw.* 18:311–19
149. Coppersmith G, Dredze M, Harman C. 2014. Quantifying mental health signals in Twitter. In *Proceedings of the Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality*, pp. 51–60. Stroudsburg, PA: Assoc. Comput. Linguist.
150. Wang Y, Weber I, Mitra P. 2016. Quantified self meets social media: sharing of weight updates on Twitter. In *Proceedings of the 6th International Conference on Digital Health Conference (DH '16)*, pp. 93–97. New York: Assoc. Comput. Mach.
151. Guntuku SC, Yaden DB, Kern ML, Ungar LH, Eichstaedt JC. 2017. Detecting depression and mental illness on social media: an integrative review. *Curr. Opin. Behav. Sci.* 18:43–49

152. Chen J, Chen H, Wu Z, Hu D, Pan JZ. 2017. Forecasting smog-related health hazard based on social media and physical sensor. *Inform. Syst.* 64:281–91
153. Oh O, Kwon K, Rao H. 2010. An exploration of social media in extreme events: rumor theory and Twitter during the Haiti earthquake 2010. In *Proceedings of the 31st International Conference on Information Systems (ICIS 2010)*, Pap. 231. https://aisel.aisnet.org/icis2010_submissions/231
154. McCullom R. 2018. A murdered teen, two million tweets and an experiment to fight gun violence. *Nature* 561:20–22
155. Lo AS, Esser MJ, Gordon KE. 2010. Youtube: a gauge of public perception and awareness surrounding epilepsy. *Epilepsy Behav.* 17:541–45
156. Betton V, Borschmann R, Docherty M, Coleman S, Brown M, Henderson C. 2015. The role of social media in reducing stigma and discrimination. *Br. J. Psychiatry* 206:443–44
157. Ladea M, Bran M, Claudiu SM. 2016. Online destigmatization of schizophrenia: a Romanian experience. *Eur. Psychiatry* 33:S276
158. Lydecker JA, Cotter EW, Palmberg AA, Simpson C, Kwitowski M, et al. 2016. Does this tweet make me look fat? A content analysis of weight stigma on Twitter. *Eat. Weight Disord.* 21:229–35
159. Witzel CT, Guise A, Nutland W, Bourne A. 2016. It starts with me: privacy concerns and stigma in the evaluation of a Facebook health promotion intervention. *Sex. Health* 13:228–33
160. Pacheco DF, Pinheiro D, Cadeiras M, Menezes R. 2017. Characterizing organ donation awareness from social media. In *2017 IEEE 33rd International Conference on Data Engineering (ICDE)*, pp. 1541–48. New York: IEEE
161. Engel J. 2003. Bringing epilepsy out of the shadows. *Neurology* 60:1412–12
162. Kerson TS. 2012. Epilepsy postings on YouTube: exercising individuals' and organizations' right to appear. *Soc. Work Health Care* 51:927–43
163. Fiest KM, Birbeck GL, Jacoby A, Jette N. 2014. Stigma in epilepsy. *Curr. Neurol. Neurosci. Rep.* 14:444
164. Sartorius N, Schulze H. 2005. *Reducing the Stigma of Mental Illness: A Report from a Global Programme of the World Psychiatric Association*. Cambridge, UK: Cambridge Univ. Press
165. Patel R, Chang T, Greysen SR, Chopra V. 2015. Social media use in chronic disease: a systematic review and novel taxonomy. *Am. J. Med.* 128:1335–50
166. McNeil K, Brna P, Gordon K. 2012. Epilepsy in the Twitter era: a need to re-tweet the way we think about seizures. *Epilepsy Behav.* 23:127–30
167. Payton FC, Kvasny L. 2016. Online HIV awareness and technology affordance benefits for black female collegians—maybe not: the case of stigma. *J. Inform. Health Biomed.* 23:1121–26
168. Silenzio VM, Duberstein PR, Tang W, Lu N, Tu X, Homan CM. 2009. Connecting the invisible dots: reaching lesbian, gay, and bisexual adolescents and young adults at risk for suicide through online social networks. *Soc. Sci. Med.* 69:469–74
169. Lupton D. 2016. The use and value of digital media for information about pregnancy and early motherhood: a focus group study. *BMC Pregnancy Childbirth* 16(1):171
170. Harpel T. 2018. Pregnant women sharing pregnancy-related information on Facebook: Web-based survey study. *J. Med. Internet Res.* 20:e115
171. Bartholomew MK, Schoppe-Sullivan SJ, Glassman M, Dush CMK, Sullivan JM. 2012. New parents' Facebook use at the transition to parenthood. *Family Relations* 61:455–69
172. Llorente A, Garcia-Herranz M, Cebrian M, Moro E. 2015. Social media fingerprints of unemployment. *PLOS ONE* 10:e0128692
173. Ramagopalan SV, Simpson A, Sammon C. 2020. Can real-world data really replace randomised clinical trials? *BMC Med.* 18:13
174. Pfeffer J, Mayer K, Morstatter F. 2018. Tampering with Twitter's sample API. *EPJ Data Sci.* 7:50
175. Ruths D, Pfeffer J. 2014. Social media for large studies of behavior. *Science* 346:1063–64
176. Jensen EA. 2017. Putting the methodological brakes on claims to measure national happiness through Twitter: methodological limitations in social media analytics. *PLOS ONE* 12:e0180080
177. Shaban-Nejad A, Michalowski M, Buckeridge DL. 2018. Health intelligence: how artificial intelligence transforms population and personalized health. *NPJ Digit. Med.* 1:53

178. Perrin A, Anderson M. 2019. Share of U.S. adults using social media, including Facebook, is mostly unchanged since 2018. *Fact Tank*, Pew Res. Cent., April 10. <https://www.pewresearch.org/fact-tank/2019/04/10/share-of-u-s-adults-using-social-media-including-facebook-is-mostly-unchanged-since-2018/>
179. Lenormand M, Picornell M, Cantú-Ros OG, Tugores A, Louail T, et al. 2014. Cross-checking different sources of mobility information. *PLOS ONE* 9:e105184
180. Prieto VM, Matos S, Álvarez M, Cacheda F, Oliveira JL. 2014. Twitter: a good place to detect health conditions. *PLOS ONE* 9:e86191
181. Luckerson V. 2015. Here's how Facebook's news feed actually works. *Time*, July 9. <https://time.com/collection-post/3950525/facebook-news-feed-algorithm/>
182. Bagrow JP, Liu X, Mitchell L. 2019. Information flow reveals prediction limits in online social activity. *Nat. Hum. Behav.* 3:122–28
183. Ginsberg J, Mohebbi MH, Patel RS, Brammer L, Smolinski MS, Brilliant L. 2009. Detecting influenza epidemics using search engine query data. *Nature* 457:1012–14
184. Yang CC, Yang H, Jiang L, Zhang M. 2012. Social media mining for drug safety signal detection. In *Proceedings of the 2012 International Workshop on Smart Health and Wellbeing*, pp. 33–40. New York: Assoc. Comput. Mach.
185. Rocha LM, Börner K, Miller WR. 2019. *myAURA: personalized web service or epilepsy management*. Grant Announc., Grantome Database, accessed Nov. 29. https://hsrproject.nlm.nih.gov/view_hsrproj_record/20191123
186. Miller WR, Gesselman AN, Garcia JR, Groves D, Buelow JM. 2017. Epilepsy-related romantic and sexual relationship problems and concerns: indications from internet message boards. *Epilepsy Behav.* 74:149–53
187. CDC (Cent. Dis. Control Prev.). 2019. *About chronic diseases*. Fact Sheet, Natl. Cent. Chron. Dis. Prev. Health Promot., Atlanta, GA, accessed Nov. 29. <https://www.cdc.gov/chronicdisease/about>
188. Grady PA, Gough LL. 2014. Self-management: a comprehensive approach to management of chronic conditions. *Am. J. Public Health* 104:e25–31
189. Cassidy JJ, Bernasek SM, Bakker R, Giri R, Peláez N, et al. 2019. Repressive gene regulation synchronizes development with cellular metabolism. *Cell* 178:980–92
190. Dolan RJ. 2002. Emotion, cognition, and behavior. *Science* 298:1191–94
191. Nofsinger JR. 2005. Social mood and financial economics. *J. Behav. Finance* 6:144–60
192. Ruiz R. 2016. Why scientists think your social media posts can help prevent suicide. *Mashable*, June 26. <https://mashable.com/2016/06/26/suicide-prevention-social-media/>
193. Margolis R, Derr L, Dunn M, Huerta M, Larkin J, et al. 2014. The National Institutes of Health's Big Data to Knowledge (BD2K) initiative: capitalizing on biomedical big data. *J. Am. Med. Inform. Assoc.* 21:957–58



Annual Review of
Biomedical Data
Science

Volume 3, 2020

Contents

Deciphering Cell Fate Decision by Integrated Single-Cell Sequencing Analysis <i>Sagar and Dominic Grün</i>	1
Knowledge-Based Biomedical Data Science <i>Tiffany J. Callahan, Ignacio J. Tripodi, Harrison Pielke-Lombardo, and Lawrence E. Hunter</i>	23
Infectious Disease Research in the Era of Big Data <i>Peter M. Kasson</i>	43
Spatial Metabolomics and Imaging Mass Spectrometry in the Age of Artificial Intelligence <i>Theodore Alexandrov</i>	61
Protein–Protein Interaction Methods and Protein Phase Separation <i>Castrense Savojardo, Pier Luigi Martelli, and Rita Casadio</i>	89
Data Integration for Immunology <i>Silvia Pineda, Daniel G. Bunis, Idit Kost, and Marina Sirota</i>	113
Computational Methods for Analysis of Large-Scale CRISPR Screens <i>Xueqiu Lin, Augustine Chemparakathy, Marie La Russa, Timothy Daley, and Lei S. Qi</i>	137
Computational Methods for Single-Particle Electron Cryomicroscopy <i>Amit Singer and Fred J. Sigworth</i>	163
Immunoinformatics: Predicting Peptide–MHC Binding <i>Morten Nielsen, Massimo Andreatta, Bjoern Peters, and Søren Buus</i>	191
Analytic and Translational Genetics <i>Konrad J. Karczewski and Alicia R. Martin</i>	217
Mobile Health Monitoring of Cardiac Status <i>Jeffrey W. Christle, Steven G. Hershman, Jessica Torres Soto, and Euan A. Ashley</i>	243
Statistical Methods in Genome-Wide Association Studies <i>Ning Sun and Hongyu Zhao</i>	265

Biomedical Data Science and Informatics Challenges to Implementing Pharmacogenomics with Electronic Health Records <i>James M. Hoffman, Allen J. Flynn, Justin E. Juskewitch, and Robert R. Freimuth</i>	289
Identifying Regulatory Elements via Deep Learning <i>Mira Barshai, Eitamar Tripto, and Yaron Orenstein</i>	315
Computational Methods for Single-Cell RNA Sequencing <i>Brian Hie, Joshua Peters, Sarah K. Nyquist, Alex K. Shalek, Bonnie Berger, and Bryan D. Bryson</i>	339
Analysis of MRI Data in Diagnostic Neuroradiology <i>Saima Rathore, Ahmed Abdulkadir, and Christos Davatzikos</i>	365
Supercomputing and Secure Cloud Infrastructures in Biology and Medicine <i>Cathrine Jespersgaard, Ali Syed, Piotr Chmura, and Peter Løngreen</i>	391
Computational Approaches for Unraveling the Effects of Variation in the Human Genome and Microbiome <i>Chengsheng Zhu, Maximilian Miller, Zishuo Zeng, Yanran Wang, Yannick Mablich, Ariel Aptekmann, and Yana Bromberg</i>	411
Mining Social Media Data for Biomedical Signals and Health-Related Behavior <i>Rion Brattig Correia, Ian B. Wood, Johan Bollen, and Luis M. Rocha</i>	433

Errata

An online log of corrections to *Annual Review of Biomedical Data Science* articles may be found at <http://www.annualreviews.org/errata/biodatasci>