How to Use Social Media Data for Political Science Research

Pablo Barberá
University of Southern California
pbarbera@usc.edu

Zachary C. Steinert-Threlkeld
University of California - Los Angeles
zst@luskin.ucla.edu

Forthcoming in Curini, L., and Franzese, R. (eds) *The SAGE Handbook of Research Methods in Political Science and International Relations*, London: Sage, 2019

Abstract

Citizens across the globe spend an increasing proportion of their daily lives on social media platforms. Their activities on the sites generate granular, time-stamped footprints of human behavior and personal interactions. This chapter offers an overview of existing research that uses social media data in the fields of Political Science and International Relations. We discuss two types of studies: those where social media is being used as a source of data to study, e.g. political networks or public opinion, and those focusing on how social media transforms different political phenomena, ranging from ideological polarization and misinformation to election campaigns. Our review also offers an in-depth analysis of the main challenges of this type of data, such as the different sources of bias that limit the generalizability of findings, the difficulty of connecting online and offline behavior, and concerns about reproducibility and ethics. To illustrate the opportunities and limitations of social media data, we also provide an applied example using Twitter as event data to study the dynamics of protest movements in Egypt and Bahrain in 2011.

Citizens across the globe spend an increasing proportion of their daily lives on social media websites, such as Twitter, Facebook or Instagram. Their activities on the sites generate granular, time-stamped footprints of human behavior and personal interactions, sometimes with longitude and latitude coordinates. A sizable proportion of these digital traces have to do with politics – social media is an increasingly popular source of political news, as well as a forum for political debates where virtually every political candidate running for office is now present.

The data generated from these interactions is in many cases freely available for research purposes and provides a depth and breadth that was unimaginable even one decade ago. Its high degree of spatial and temporal granularity allows the study of behavior at low levels of aggregation but also at a more macro scale and from a comparative perspective. The fact that human behavior is observed unobtrusively also facilitates collecting data at a larger scale and reduces certain types of biases. This set of advantages makes social media data a new and exciting source of information to study key questions about political and social behavior.

At the same time, the volume and heterogeneity of this new type of data present unprecedented methodological challenges. Several sources of bias can limit the generalizability of our findings. Often it is difficult to connect online interactions with offline behavior. Unlike data created by governments or collected by researchers, the data-generating process is not always known. For example, we don't know if the platform was running a randomized experiment at the time of data collection or whether content was being blocked by the internet provider. And, of course, many scholars have raised concerns about the ethics of collecting data from individuals without in some cases obtaining their informed consent.

In this chapter, we provide a fair assessment of these advantages and limitations in the use of social media as data generators for research in the social sciences, with a particular emphasis on the study of political behavior. We illustrate these strengths and weaknesses with examples from two types of studies: those that where social media is being used merely as a source of data, and those where the focus is on how social media is transforming different political phenomena. The first group includes research that uses social media to measure public opinion, the ideology of citizens and elites, the structure of social networks, government censorship, conflict dynamics, and elite rhetoric; and where social media sites are used as a new space to conduct affordable field experiments. The second set of studies deals with questions such as how social media platforms contribute to the success of collective action events, how they are transforming election campaigns, and whether their usage is contributing to greater political polarization and the spread of misinformation.

Throughout the chapter, we consider "social media data" as any type of information obtained from websites whose value primarily comes through user interactions. The most well-known examples include Facebook, Instagram, Twitter, and Sina Weibo, all of which provide infrastructure to facilitate users sharing, and commenting on, content. The content can be user generated, such as status updates or photos from personal devices, or created by a third-party, such as a newspaper, and shared by users. Excluded from this definition are sites such as reddit and YouTube, whose primary purpose is to surface content from elsewhere and comment on it. Despite that stipulation, the issues described throughout this chapter apply almost equally to these social media-like websites.

To further explore the opportunities and limitations of social media data, we also provide an in-depth description of an applied example that uses Twitter data to study the dynamics of protest movements in Egypt and Bahrain in 2011.

Advantages of social media data

Perhaps the most important advantage of social media compared to traditional sources of data in the social sciences, such as surveys or government records, is the ability to unobtrusively collect information about a large sample of individuals with only minimal costs. Being able to observe subjects in a real-world environment where they are engaging in ordinary interactions reduces the likelihood that individuals change their behavior in response to knowing they're part of a social science research project, which minimizes Hawthorne and social desirability biases.[1] Although the set of APIs available for academic research has shrunk in recent years (Freelon, 2018), large amounts of such data are still freely available at a minimal cost, which facilitates comparative and longitudinal analysis of any phenomenon.

This new ability to collect information about human behavior at scale is revolutionizing the state of the art across many fields. One such example is the study of social networks. This literature had traditionally relied on either small-scale networks, such as those related to high-school classrooms (Moreno, 1934) or social groups (Zachary, 1977), or on partial views of a network reconstructed through survey responses to questions about social ties, such as in the study of communication networks (Huckfeldt, Johnson and Sprague, 2004; Mutz, 2002). Large-scale network datasets from social media sites have offered new evidence to answer some of they key standing questions within this field (see e.g. Bakshy, Messing and Adamic, 2015; Ugander et al., 2011; Lerman, Ghosh and Surachawala, 2012).

A second notable advantage of social media data is that its homogeneity in format and content facilitates systematic comparisons for different types of ac-

[1] Of course, the fact that behavior on social media takes place in public may introduce other types of social desirability bias.

tors, across multiple countries, and over time. This homogeneity applies to multiple dimensions: constrains on text length, the use of similar language and textual marks (e.g. hashtags), the ability of different types of actors to interact directly (e.g. via shares or retweets), among other many platform affordances and features. This strength of social media data is already leading to important breakthroughs in the analysis of political mobilization (Anastasopoulos and Williams, 2017; Freelon, McIlwain and Clark, 2016) and agenda-setting dynamics (Barberá et al., 2014), as well as the comparative study of political polarization (Bright, 2018) and censorship (Hobbs and Roberts, 2018).

Third, social media data offers unparalleled temporal and spatial granularity. It allows researchers to observe longitudinal trends to identify events and also to study geographic patterns. This advantage can be especially relevant in autocracies or conflict areas, since it would be difficult or impossible to gather real-time data in such situations. Social media data have proven particularly useful for understanding the dynamics of Syria's protests and subsequent civil war (O'Callaghan et al., 2014; Freelon, Lynch and Aday, 2015; Kostyuk and Zhukov, 2017); the same is true for Ukraine (Gruzd and Tsyganova, 2015; Zhukov, 2015; Wilson, 2017; Driscoll and Steinert-Threlkeld, 2018). In the last section of this chapter we offer an additional example that focuses on protests in Egypt and Bahrain leading to the Arab Spring. For technical discussions of using temporal data to detect events, see the computer science literature on event detection (Sakaki, Okazaki and Matsuo, 2010; Cadena et al., 2015; Alsaedi, Burnap and Rana, 2017) or work by Golder and Macy (2011) on the relationship between diurnal and seasonal mood patterns. For the possibility of social media data for the generation of political event data, see Zhang and Pan (2019) and Steinert-Threlkeld (2019).

Fourth, despite obvious concerns about representativeness of samples obtained from social media (a point to which we return next), for some groups of actors virtually the entire population is present on social networking platforms. For example, over 85% of world leaders have active Twitter or Facebook accounts (Barberá and Zeitzoff, 2017) and virtually all Members of the U.S. Congress also maintain profiles on these sites (Pew Research Center, 2017), making it possible to make externally valid inferences about their communication strategies and rhetoric with data obtained from these accounts.

As we consider samples of ordinary citizens, it is worth emphasizing that representativeness of social media users is only a concern if the behavior under study is thought to vary according to variables by which users on social media differ from the population from which they are drawn. For example, Malik et al. (2015) find that a 1% increase in the population size of a census block group correlates with 5.7% fewer geolocated Twitter users. As a result, using Twitter to understand behaviors correlated with living in rural areas, such as sentiments towards a trade war or support for political candidates, is likely to underestimate those behaviors in the population.[2] Similarly for studying older or poorer Americans.

On the other hand, many behaviors should be less sensitive to non-representative samples. For example, factors affecting information contagion – weak and strong ties, the influence of media and celebrities, etc. – probably do not vary by age, location, or race. During the Arab Spring, users of social media were certainly not representative of their populations (Ghannam, 2011; Breuer, 2012); the vast majority of protesters in Tahrir Square, for example, learned about the protests from satellite television or face-to-face communication (Wilson, Gosling and Graham, 2012). It is not clear, however, why patterns of information diffusion found on Twitter at this time (Brym et al., 2014; Borge-Holthoefer et al., 2015; Steinert-Threlkeld, 2017*b*) should be assumed to be distinct from information spreading

[2] Observing a group over time obviates the numerator problem.

outside of social media. Other outcomes, such as diurnal patterns of behavior or user to user interactions, for example, should also not vary on observables.[3]

The fifth advantage of social media data, that many well-known offline behaviors also occur online, supports the claim – representativeness issues matter but probably not as much as feared – of the previous paragraph. Social media users' ideology is recoverable from the structure of their social network and maps onto offline estimates of ideology (Barbera, 2015; Bond and Messing, 2015). Dunbar's Number (Dunbar et al., 2015), attitudinal homophily (Bliss et al., 2012), diurnal activity (Golder and Macy, 2011), and geographic constraints (Takhteyev, Gruzd and Wellman, 2012) occur offline and online, and Facebook users' scores on the Big Five personality traits (extraversion, neuroticism, conscientiousness, agreeableness, and openness to experience) are recoverable from their behavior on the site (Gosling et al., 2011). Twitter users also exhibit homophily with respect to age, gender, and political affiliation (Zamal, Liu and Ruths, 2012), and homophily of likes on Facebook permits similar behavioral and personality inference (Kosinski, Stillwell and Graepel, 2013*a*). The preponderance of evidence suggests it is reasonable to expect that many, perhaps most, behaviors studied using social media are analogues to what would be observed offline, if it were possible to observe these behaviors at scale offline. It is rarely possible, of course, to observe these behaviors at scale offline.

Limitations of social media data

A fair assessment of how social media data can be used in social science research requires also a discussion of its key limitations, as well as potential ways to address them. Probably the most important challenge to overcome, as we discussed

[3] "On observables" is, of course, a major qualification.

in the previous section, is that social media users are not a representative sample of the population in any given country. As a result, an analysis of social media data without any type of adjustment may reveal patterns that are not generalizable. Despite this valid concern, it is worth making two points. First, samples that are not representative can still be scientifically relevant. For example, we may treat the sample of users posting about politics on Twitter as a set of "opinion leaders" that can be more influential than other ordinary citizens. And second, there are actually different types of sampling bias, and it is important to quantify them. We could have sampling bias whenever sociodemographic characteristics are correlated with *both* our outcome variables and the propensity to be present on social media. But there may also be self-selection *within* samples of social media users, particularly when samples are collected at the tweet level, based on whether it mentions or not a set of keywords. For example, if our sample includes tweets mentioning names of parties or political candidates, it will oversample individuals with extreme political identities, because they tend to tweet about politics more frequently (Barberá and Rivero, 2015).[4]

One potential solution to this set of problems would be to apply similar weighting methods as in survey research, where low response rates can lead to similar concerns about sampling bias. However, we are limited by what Golder and Macy (2014) call the "privacy paradox": "[social media] data are at once too revealing in terms of privacy protection, yet also not revealing enough in terms of providing the demographic background information needed by social scientists".

A different type of sampling issues that is also worth mentioning are those due to black-box proprietary sampling algorithms used by social media companies. For example, Morstatter et al. (2013) showed that the Twitter API does

[4]Keyword self-selection can be mitigated by connecting to the streaming API and downloading a 1% sample of tweets in real time. A downside of this approach is that events or keywords that are not very popular are less likely to appear in the stream since it is a sample.

not return a truly random sample if we are collecting from the "spritzer" and not sampling using keywords.[5] Why this difference occurs, however, is unclear from existing work, as Twitter populates the spritzer based on a tweet's millisecond timestamp (Pfeffer and Mayer, 2018). More work is needed, however, to determine how sensitive this result is to studying particular topics or times, as well as to whether the observed differences are relevant enough to be significant.

A different type of challenge lies on the ability to connect online and offline behavior. The fact that we observe a pattern in an online setting does not mean that it would necessarily replicate in an offline setting. This could be due to affordances of the platforms – e.g. the fact that anonymity is easier in online settings makes vitriol and incivility more likely to occur in interpersonal communications – or to the type of ties that people develop on social media. However, as discussed in the previous section, there's increasing evidence that this is not the case. For example, two recent studies demonstrate that the structure of networks and the nature of social ties are similar in social media and offline networks (Bisbee and Larson, 2017; Jones et al., 2013).

An additional difficulty of conducting network analysis using social media data, especially Twitter, is defining what an *edge* is. Connections on social media form more easily than offline, so it is common to infer edge strength via other behavior. On Facebook, appearing together in a photograph or tagging someone in a status is commonly used to define "true" friends. On Twitter, researchers will infer an edge if a user retweets or mentions another user; for topic modeling, edges usually come from the co-occurrence of hashtags. For Twitter, retweet and user mentions are preferred because they can be easily extracted from the streaming API, whereas finding an account's followers requires more

[5] Morstatter et al. (2013) show divergence in trends in (1) the number of tweets, (2) topics, and (3) some network features in tweets related to Syria from 12.14.2011-01.10.2012. Overall, it is clear that a random sample collected from streaming API may not be a perfectly representative sample of all of Twitter.

work and is severely hindered by rate limits Twitter imposes on how frequently a researcher can request data. A shortcoming of this approach is that it usually creates cross-section data, removing any temporal information about a network. Steinert-Threlkeld (2017*a*) introduces a method that works within the Twitter rate limits to measure accounts' influence as it changes every day. This approach requires more engineering than relying on retweets and user mentions delivered via the REST API, but it permits longitudinal observation.

A third limitation of social media data is replicability, for three reasons. First, sharing raw data is difficult. Twitter, the most common social media platform studied, only allows an individual to share 50,000 tweets per day "via non-automated means"; this quantity is not large.[6] For studies with millions of tweets, sharing tweets therefore requires either moderate amounts of programming to build a front-end to the dataset to ensure replicators do not receive more than 50,000 tweets per day or a human in the loop for however many days are required to share the entire dataset. Neither option is easy, and both detract from research productivity.[7] The best method for sharing a full dataset is therefore to share the identification number of each tweet, as Twitter has no cap on that quantity *when used for academic research*.[8]. For Python and R scripts to download tweets using their ID numbers, see Steinert-Threlkeld (2018).

Sharing tweet IDs invokes the second reason, which is programmer skill. Because downloading tweets requires connecting to Twitter's application programming interface (API), and the returned data is made available in a format (JSON) with which many researchers may not be familiar, replication requires

[6]Our discussion here focuses on Twitter, as the main social media platform offering easy access to public data for research purposes. Although similar data used to be available for Instagram and Facebook, recent changes to their platform policies mean that their APIs are essentially no longer available for research purposes (Freelon, 2018).

[7]An alternative would be to develop a framework that can produce synthetic data with similar properties as the full dataset while ensuring privacy (Raab, Nowok and Dibben, 2016), but to our knowledge this approach has not been developed for social media datasets yet.

[8]If the ID numbers are distributed for non-academic purposes, no more than 1.5 million per 30 day period can be shared.

programming skills that many replicators may not possess. A new R package, `rehydratoR`, simplifies this process. It takes a list of tweet IDs and saves the results to .csv files, in addition to accounting for Twitter rate limits (Coakley, Kirkpatrick and Steinert-Threlkeld, 2019).

The third reason is "post rot." If a user deletes a tweet or account, that tweet (and all tweets from that account) is not later recoverable. A tweet included in a study may therefore not be available to a later replicator. While the rate of tweet or user rot is not known, as far as we are aware, anecdotal reports suggest it is about $\frac{1}{10}$ of tweets after one year. Whether or not this rot changes a replicators' inference is not known, although there is some evidence that rot does not occur randomly (Timoneda, 2018). Even if the overall findings may not change, point estimates probably will.

If replicators start with the aggregated data that authors use for their analysis, then these three reasons are obviated. Restrictions from Institutional Review Boards also often require that authors do not share tweet IDs because it would allow other researchers to identify the subjects in the study. In that case, sharing aggregated data is the only approach to replication. This outcome is especially the case in experiments on social media (Munger, 2016; Siegel and Badaan, 2018).

Ethical concerns are, of course, important beyond the discussion about replicability. A controversial recent study that experimentally manipulated the content of Facebook users' news feed to explore whether emotions are contagious (Kramer, Guillory and Hancock, 2014) opened a debate about what the notion of informed concern means in an environment in which companies constantly manipulate features of their product using experimental methods. Even observational studies with publicly-available data can raise ethical concerns because individuals may not have the expectation that their data is being analyzed for research. Even if data is de-anonymized, often it is possible to re-identify personal data (Zimmer, 2010). It is important to note, however, that this problem is present in any other study that uses individual-level data, including those that use survey data. Some recent work in the area of differential privacy (Dwork, 2008) offers exciting new possibilities to strike a balance between data access and privacy.

The fifth limitation is norms of users and companies. On user norms, the primary difference between Facebook and Twitter, and one that is not often appreciated, is that users on Twitter maintain a norm of public production while most users on Facebook opt for privacy. A small minority of Twitter users maintain protected profiles, making their data as inaccessible as the vast majority of Facebook users'. Companies also maintain their own norms around how much data to expose via their APIs, so the kind of research possible is at the whim of companies' API policies. For example, Facebook and Instagram used to provide much more user and graph information via their APIs, but they have become much more restrictive in response to recent controversies about data protection. Twitter has made it more difficult to obtain a developer account, the first step to collecting data, though changes to its API have otherwise been positive.

Using social media to measure political behavior

To illustrate the strengths and limitations of the use of social media data in the social sciences, we now offer an overview of recent research within this subfield. Some of the earliest work that took advantage of the opportunities that this new source of data offers was focused on measuring public opinion. Initial claims that lauded the potential of tweets to predict election results (Tumasjan et al., 2010) were soon rebuked by more systematic analysis demonstrating that those predictions were in many instances no better than random guesses (Gayo Avello, Metaxas and Mustafaraj, 2011). With some notable exceptions, such as the work of Ceron et al. (2014), the probability that we will be replacing public opinion surveys with metrics based on social media seems still unlikely (Klašnja et al., 2016). However, there is clear value in social media as a *complement* to survey data, both as an early indicator of changes in public opinion and a possible signal on unpolled topics or areas (Beauchamp, 2017).

Even if aggregate-level opinion is hard to measure, there is plenty of evidence that social media can reveal a lot of information about citizens' characteristics and behavior. Kosinski, Stillwell and Graepel (2013*b*) and Youyou, Kosinski and Stillwell (2015) found that Facebook likes are highly predictive of private traits such as party preference, age, gender, sexual orientation, alcohol use, psychological traits, etc. It was this type of analysis that Cambridge Analytica allegedly used to target voters with personalized ads through social media; although there is no evidence regarding its effectiveness. Regarding the political domain more specifically, Barbera (2015) and Bond and Messing (2015) demonstrated that the political accounts that each citizen likes on Facebook or follows on Twitter offer enough data to estimate political ideology on a continuous scale (liberal to conservative) with high accuracy. Other recent work (Radford and Sinclair, 2016) suggests that the text of social media messages could also be used to estimate political preferences.

Since most autocratic governments around the world try to exercise some type of control over social media sites, we can also use data from these sources to better understand digital repression strategies. For example, King, Pan and Roberts (2013, 2014) found that it is possible to "reverse engineer" online censorship using a sophisticated system that scrapes social media in real-time and then checks what has been deleted and why. Their analysis revealed that the Chinese government prioritizes deleting content that could lead to collective action but that they allow criticism of the government. Another way governments can exercise digital repression is through trolls or bots that flood a platform with pro-regime information or with topics to distract from a politically sensitive issue. By muddying the waters, this tactic allows the government to silence dissidents' attempts at political coordination while maintaining plausible deniability that censorship has occurred (Little, 2015; King, Pan and Roberts, 2017; Keller et al., 2017). These tactics also make it difficult for researchers to separate true beliefs from pro-government noise. One potential solution is to identify bots using bot detection methods (Ferrara et al., 2016), although this becomes more difficult as we consider trolls, which are accounts controlled by humans.

In the same way that press releases or campaign speeches are often used to examine elite rhetoric (see e.g. Grimmer, Messing and Westwood, 2012), tweets and Facebook posts by politicians can also be a useful source of data to understand political communication. The series of reports on congressional rhetoric released by the Pew Research Center are a good example. There is also a growing body of comparative work on topics such as campaign strategies or populism (Theocharis et al., 2016; Nulty et al., 2016; Stier et al., 2017), including for instance politicians' attempts to overwhelm social media with false messages (a technique also known as *astroturfing*) so that challengers' support appears smaller than it may actually be (Munger et al., 2018). It is also possible to recover politicians' ideology from the images they share on Facebook, which may be especially useful for estimating policy positions of challengers who do not have a voting record or sizable donations (Joo, Anastasopoulos and Steinert-Threlkeld, 2019).

At a more aggregate level, social media can reveal conflict dynamics with detail unavailable via other approaches. For example, students of repression and dissent generally code repression as a binary or categorical variable. Advances in computer vision techniques can be applied to thousands of protest photos to generate a continuous measure of state and protester violence. These techniques can also generate estimates of the race and gender of participants, as well as the number of participants (Joo, Won and Steinert-Threlkeld, 2018). Though social media tend to have some sort of location bias, it is reasonable to expect that this bias in less than in newspapers because of the lower barriers to entry and unlimited publication space (Steinert-Threlkeld and Joo, 2019). Social media data can also reveal daily changes in interstate conflict dynamics (Zeitzoff, 2011; Zeitzoff, Kelly and Lotan, 2015; Zeitzoff, 2016) and civil wars (Zhukov, 2015; Driscoll and Steinert-Threlkeld, 2018) as well as document an order of magnitude more protests in China than newspapers have (Zhang and Pan, 2019). It also has been used to study Black Lives Matter protests in the United States (Anastasopoulos and Williams, 2017; Hsuan et al., 2017).

Finally, social media can also be a space to conduct affordable field experiments. In a pioneering article, Munger (2016) developed a way to deploy experiments on Twitter using bots to convey treatments. His study discovered that, when users who use racist slurs are exposed to a bot that criticizes their actions, they change their behavior, but only when the "punishment" is coming from someone similar to them. Compared to field experiments deployed in an offline environment, administering treatments has a much lower cost, although it requires an effort in terms of hardware and programming skills. Building a bot to administer a treatment is more challenging than passively collecting data from an API endpoint. Despite this limitation, this approach represents a promising blueprint for future experimental studies with high external validity (Siegel and Badaan, 2018; Coppock, Guess and Ternovski, 2016).

Understanding how social media affects political behavior

An important part of the ongoing research that uses social media data focuses not necessarily on using these sites as a source of information about behavior, but instead on their potential a transformative political force whose effects we are only starting to understand. There's perhaps no better example of this type of work than the efforts to understand how digital technologies were a catalyst in the recent wave of protests around the world, starting with the Arab Spring. Some of the earlier optimism about the democratizing power of social media (see e.g. Valenzuela, 2013) was followed by a wave of skepticism. Authors such as Gladwell (2010) or Morozov (2012) warned against the rise of "slacktivism" fueled by social media sites, which facilitate engagement in protests but also disincentivize commitment and the type of dedicated and trained activism that can turn the revolutionary fever into action. However, recent empirical research in political science has demonstrated that it is precisely the ability of social media to bring in peripheral individuals without resources or deep engagement that can lead to the success of collective action event (Barberá, Wang, Bonneau, Jost, Nagler, Tucker and González-Bailón, 2015; Steinert-Threlkeld, 2017*b*).

This evolving narrative about the impact of technology on protest mirrors the public discussion about how social media has been used in political campaigns. Barack Obama's election campaigns are widely acknowledged to have brought the data revolution to politics (Kreiss, 2012). An important part of his success lied on understanding the native digital audience and on deploying an extensive data collection and analysis platform. In contrast, Mitt Romney's campaign famously required 22 people to approve every single tweet from the candidate's account (Kreiss, 2016). But the optimism about how digital technologies can empower grassroots movements was put to test after Donald Trump's election in 2016, and the alleged use by his campaign of micro-targeted ads that were tailored to voters' preferences. Although there is still a significant lack of research about whether targeted advertising on social media can be effective at mobilizing or persuading voters, the existing work generally finds null or very small effects, given us reasons to be skeptic (Broockman and Green, 2014; Eckles, Gordon and Johnson, 2018; Kalla and Broockman, 2018; Bond et al., 2012)

Another key research question in this field has been whether news consumption and political conversations through social media may be one of the factors explaining the recent rise in political polarization and extremism around the world. A common argument in the literature is that social networking sites make it easier for citizens to isolate themselves into communities of like-minded individuals, where political agreement is the norm and individuals can avoid being exposed to any opinion that may challenge their ideological views (Sunstein, 2018). This process could be exacerbated by ranking algorithms that filter out any content that users may dislike (Pariser, 2011). These concerns are not new – Putnam (2000) already expressed concerns about cyber-balkanization nearly two decades ago – but they have re-emerged and are often identified as an additional factor explaining the recent rise in in political polarization. Despite these concerns, empirical evidence that this process is happening is scarce: individuals are exposed to more diverse views on social media than in offline settings (Bakshy, Messing and Adamic, 2015; Fletcher and Nielsen, 2018; Barnidge, 2017), cross-ideological interactions are frequent even in relation to highly contentious topics (Barberá, Jost, Nagler, Tucker and Bonneau, 2015), and the increase in polarization in the US has been largest among those who are least likely to be active on social media (Boxell, Gentzkow and Shapiro, 2017).

Finally, perhaps the most timely topic that has received the broadest media attention is the extent to which social media contributes to the spread of misinformation. Contrary to the conventional wisdom that exposure to false news during the 2016 presidential election, a growing consensus is emerging in the literature that points to important asymmetries in the extent to which this type of information is shared and consumed. Conservative, older citizens are more likely to consume misinformation on social media (Guess, Nagler and Tucker, 2019) and exposure appears to be concentrated on a small minority of users. Whether or not being exposed to false news can actually affect attitudes and behavior is unclear, and evidence is mixed (Jamieson, 2018; Allcott and Gentzkow, 2017).

Case study: analyzing the dynamics of protest movements using Twitter data

In this last section, we demonstrate the ability of social media as a data generator through a case study of activists and government accounts from Egypt and Bahrain in early 2011, during the Arab Spring. An analysis of data collected from Twitter explores the question of whether activists can activate offline protest through their online activity on social media.

The phrase "Arab Spring" refers to the large-scale protests that occurred throughout the Middle East and North Africa from December 2010 through the end of 2011. On December 17th, 2010, Mohamed Bouazizi immolated to protest the seizure of his fruit cart. His action inspired local protests that became national, and the national protests spread to other countries once President Ben Ali fled to Saudi Arabia on January 14th, 2011. Hosni Mubarak would abdicate on February 14th, and almost every country in the region experienced some form of protest.

Though there is much evidence that the protests were spontaneous (Wilson, Gosling and Graham, 2012; Steinert-Threlkeld, 2017*b*), activists played important roles at various points. Before the protests, they held workshops about peaceful resistance, showed movies about protest or critical of a regime, and taught how to create political graffiti. These activities demonstrated to others that there exists discontent with a regime, and they do so in a way so that others know that others know there is discontent. Generating political coordination has the same affect as advertising: the more people know about a product (the protest), the more likely are people to buy it (join the protest) (Chwe, 1998). For example, Egyptian activists deliberately discussed (advertised) the January 25 protests in taxi cabs because they knew the drivers would spread the information to their passengers (Lim, 2012). However, there is still little systematic research on the extent to which activists contributed to the spread of protest – that is the research question we try to address here.

The Arab Spring presents an ideal case to demonstrate the strengths of social media data. Because the involved countries were – and largely still are – repressive, identifying networks of activists, gaining their trust, and administering surveys was (is) a resource-intensive process. For example, even after Hosni Mubarak abdicated, individuals in Tahrir Square were afraid to response to Western researchers' surveys (Tufekci and Wilson, 2012). Any work coming from that process would be difficult to compare to other situations because the researcher will have invested so much time in gathering data for a particular set of actors in one country. Even within a particular country, it would be difficult to study many movements at once.

Once a researcher has established him or herself in a repressive country, administering surveys presents additional challenges. Respondents are often suspicious of foreigners conducting research or do not want to be seen associated with them. Local enumerators may be hired, adding time and expense to the data gathering process. If an unforeseen event occurs, such as a protest, the researcher is unlikely to be able to pivot to study it. Even if the researcher can study this unexpected event, she or he will be hard pressed to gather many observations per day, not to mention many observations across many locales per day. Survey work is hard, especially in places where it is infrequently conducted.

For the reasons provided earlier in this chapter, social media data ameliorates many of these concerns. To demonstrate how they do so, we analyze 19 activists across four social movements in Egypt and Bahrain, and examine whether their social media behavior led to an increase in offline protest.

In Egypt, we focus on the April 6th movement, which started in early 2008 as a Facebook page rallying support for striking textile workers at a government enterprise in Mahalla al-Kubra, a city of 535,000 inhabitants located 70 miles north of Cairo. Large-scale strikes and protests focused on working conditions and pay had occurred since 2006, sparking a periodic series of worker actions throughout Egypt over the next two years (Beinin, 2009). The most important event was a large strike that was called for April 6, 2008, which the government reacted to by preemptively arresting activists and closing off public spaces nationwide (Gunning and Baron, 2013, pgs. 59-61). The movement persisted at a subdued level of activity – not for lack of trying – for the next three years and would become a central actor in the 2011 mobilization.

The second social movement in Egypt is the anti-sexual harassment movement. Egyptian public spaces have long been dangerous for women (Amar, 2011). As protests increased in Egypt throughout the first decade of the new millennium, so did reports of sexual assault at these events; in many cases, these assaults are linked to civilians the Interior Ministry hired for that purpose (Langohr, 2013). In response to these events, an assortment of civil society organizations emerged, most notably the Nadeem Center, Egyptian Center for Women's

Rights, and the Nazra for Feminist Studies.

We also study the two main organizations in support of human rights in Bahrain, the Bahrain Center for Human Rights (BCHR) and Bahrain Human Rights Society (BHRS), which were founded in 2002 and were still active at the time of the 2011 protests.

To study these movements, we rely on data that contains a complete history of the 19 most important activists within these movements on Twitter. The data was purchased from Sifter, a third-party reseller, and spans the period from January 11th, 2011 through April 5th, 2011. This time period was selected because it encompasses the time leading to each country's main protest period, the time during the main protest period (January 25 – February 11 in Egypt, February 14 – March 17 in Bahrain), and time after the protests. [9] Sifter returned 58,376 tweets; each includes metadata on the number of followers of the account, number of people the account follows, and a character string describing the device from which the tweet was created. Table 1 details these data.[10]

Did the activities on Twitter by these key leaders have any effect on subsequent protest events in these two countries? Contrary to the conventional wisdom that highly visible activists were determinant in the spread of protest across these two countries, here we show that there is little evidence supporting that assertion. Table 2 presents regression where we test that hypothesis through a multivariate regression on the number of protests with a range of variables related to the social media activity of leaders before those protests took place. The data on protests was extracted from the Integrated Conflict Early Warning Sys-

[9]The end of the main protest period in Egypt is defined as Mubarak's resignation. In Bahrain, protests ended following a major assault on the Pearl Roundabout, the main protest site in Manama; this assault occurred three days after Gulf Cooperation Council forces, led by Saudi Arabia, marched into Bahrain, and the Pearl Roundabout was dismantled on March 18. Protesters would not again succeed in occupying the circle.

[10]For additional information about data formats and code to collect and analyze Twitter data, see the materials available in Steinert-Threlkeld (2018).

tem (ICEWS) (Boschee et al., 2015).[11] To measure activists' social media behavior, we use a set of different variables that include the percent of tweets from activists, the percent of tweets that are activist retweets, the percent of tweets that are activists mentioning others, and the percent of tweets with hashtags that are from activists, and the percent of tweets with links from activists.[12]

[11]ICEWS is a Department of Defense project, led by Lockheed Martin and Michael Ward, that reads newspapers and extracts events. It represents a substantial modification and extension of Philip Schrodt's Kansas Events Data System (KEDS) and Textual Analysis by Augmented Replacement Instructions (TABARI) (Schrodt, Davis and Weddle, 1994; Gerner et al., 2002). ICEWS reads thousands of news sources, including non-English ones, and applies a heavily modified version of TABARI, leading to much lower rates of false positives than other machine-coded events data.

[12]Note that *Activist Hashtag* $\%_{i,t-1}$ is calculated slightly differently from *Activist Tweet* $\%_{i,t-1}$, *Activist Retweet* $\%_{i,t-1}$, *Activist Mention* $\%_{i,t-1}$, and *Activist Link* $\%_{i,t-1}$. *Activist Hashtag* $\%_{i,t-1}$ is calculated as the percent of all tweets with hashtags that are tweets from activists, but the other four take the total number of tweets from the activists' country on that day as the denominator. The variables are modeled differently to reflect the information consumption process on Twitter. When one sees a tweet on Twitter, it is presented as part of a sequence of reverse chronological tweets. If one views tweets containing a hashtag, however, all tweets in the subsequent reverse chronological sequence contain that hashtag. The determinant of the length of the latter sequence is therefore all tweets containing that hashtag while the length of all tweets one sees is better approximated by all tweets on that day.

Table 1: Descriptive Statistics from Sifter Data

Account	Followers	Friends	Tweets	Twitter.com	HTTPs	iPhone	Android	BlackBerry	Windows	Nokia	Retweets	Mention	Hashtag	Group
Shabab6april	4229	78	833	0.67	0.00	0.00	0.00	0.00	0.00	0.00	0.09	0.16	0.87	April 6
mrmeit	2377	244	16262	0.84	0.00	0.00	0.00	0.00	0.00	0.00	0.16	0.59	0.32	April 6
AsmaaMahfouz	1418	88	201	0.98	0.00	0.00	0.00	0.00	0.00	0.00	0.07	0.57	0.16	April 6
waleedrashed	694	329	3265	0.77	0.00	0.00	0.00	0.22	0.00	0.00	0.36	0.76	0.24	April 6
Seldemerdash	1616	277	11832	0.14	0.00	0.67	0.00	0.00	0.00	0.00	0.19	0.54	0.22	Anti-SH
Anti-SHmap	726	154	234	0.51	0.00	0.11	0.00	0.00	0.00	0.00	0.32	0.43	0.38	Anti-SH
SorayaBahgat	315	267	344	0.11	0.00	0.00	0.00	0.89	0.00	0.00	0.15	0.66	0.42	Anti-SH
MariamKirollos	127	37	564	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.52	0.66	0.60	Anti-SH
Ribeska	15	53	23	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.42	0.33	0.83	Anti-SH
ZeinabSabet	2	7	5	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	Anti-SH
alaa	13109	371	38755	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.28	0.92	0.29	NoMilTrials
Monasosh	8595	251	9479	0.47	0.00	0.00	0.00	0.53	0.00	0.00	0.18	0.68	0.40	NoMilTrials
NABEELRAJAB	7285	433	2985	0.80	0.00	0.16	0.00	0.00	0.00	0.00	0.10	0.22	0.84	Hum. Rights
BahrainRights	6687	574	4641	0.00	0.00	0.63	0.00	0.00	0.00	0.00	0.28	0.68	0.92	Hum. Rights
MARYAMALKHAWAJA	4790	187	1112	0.45	0.00	0.30	0.00	0.25	0.00	0.00	0.43	0.54	0.76	Hum. Rights
angryarabiya	3274	144	2575	0.80	0.00	0.00	0.00	0.13	0.00	0.00	0.10	0.62	0.28	Hum. Rights
SAIDYOUSIF	765	56	1574	0.49	0.00	0.51	0.00	0.00	0.00	0.00	0.34	0.47	0.84	Hum. Rights

NB: Accounts are sorted alphabetically within each movement. Columns from Twitter.com through Nokia refer to the percent of tweets from each source; for example, 84% of @mrmeit's tweets are through Twitter's website. Columns Retweets through Hashtag refer to the percent of each account's tweets which are a retweet, contain a user mention, or contain at least one hashtag.

Our first model in Table 2 shows that none of the variables related with activists' behavior on Twitter predicts subsequent protest. And in fact, *Activist Coordination Tweet* $\%_{i,t-1}$, the percent of tweets coordinating protests that are from activists, negatively correlates with protest.[13]. The second model in Table 2 uses a different measure of activist coordination but is otherwise identical to the first. Here, *Activist Coordination*$_{i,t-1}$ is the interaction of Non-Activist Coordination$_{i,t-1}$ and Activist Hashtag $\%_{i,t-1}$. Non-Activist Coordination$_{i,t-1}$ is the Gini coefficient for hashtags per country-day; if a large percentage of those hashtags are produced by activists, then hashtag coordination comes from activists. This measure is not statistically significant, while *Activist Coordination Tweet* $\%_{i,t-1}$ remains so. Note that in both models Non-Activist Coordination$_{i,t-1}$ is significant with a p-value less than .01.

Activists in both countries frequently use mobile devices for communication, but they also used desktop computers. One may argue that our results in the previous set of regression models are null because they do not focus specifically on calls for actions from whenever the activists are actually in the streets engaging in political protest. To address this concern, we take advantage of a key piece of metadata that comes with tweets, the "tweet_source" field. This field is a string created by Twitter to reflect the provenance of a tweet. We find that our results do not change when we analyze specifically the proportion of tweets originated on a mobile device. Although the percent of an activist's tweets that are from a mobile phone does positively correlate with subsequent protest and has a p-value between .05 and .10, the result does not hold when that value is interacted with the percent of tweets that are about coordination. This result makes sense, as a tweet does not say whether or not it comes from a mobile device, so a tweet from a mobile device does not provide a signal to others that the author has mo-

[13]These coordination tweets are those determined by a supervised topic model. For more details, see Steinert-Threlkeld (2017*b*)

Table 2: Main Results

	DV: $\text{Protest}_{i,t}$	
	(1)	(2)
Activist Tweet $\%_{i,t-1}$	2.58	1.72
	(6.98)	(7.47)
Activist Retweet $\%_{i,t-1}$	5.65	6.26
	(5.35)	(5.54)
Activist Mention $\%_{i,t-1}$	-2.23	-1.15
	(7.32)	(8.08)
Activist Hashtag $\%_{i,t-1}$	-1.22	-1.27
	(1.13)	(1.14)
Activist Link $\%_{i,t-1}$	9.75	11.19
	(11.80)	(12.18)
$\text{Protest}_{i,t-1}$	0.01	0.01
	(0.01)	(0.01)
$\text{Repression}_{i,t-1}$	0.01	0.01
	(0.02)	(0.02)
Non-Activist $\text{Coordination}_{i,t-1}$	5.31^{***}	5.33^{***}
	(0.85)	(0.85)
Activist Coordination Tweet $\%_{i,t-1}$	-253.02^{**}	-218.72^{*}
	(116.86)	(132.53)
Activist $\text{Coordination}_{i,t-1}$		-181.11
		(393.02)
Intercept	-2.06^{***}	-2.08^{***}
	(0.49)	(0.50)
Country FE	Yes	Yes
N	168	168
Log Likelihood	-473.42	-473.33

$^{*}p < .1$; $^{**}p < .05$; $^{***}p < .01$
NB: All activist variables except for Activist Hashtag $\%_{i,t-1}$ use the total tweets from country_i as the denominator. Activist $\text{Hashtag}\%_{i,t-1}$ uses the total tweets with hashtags from country_i.

bilized.

Table 3: Robustness Checks - Tweets from Phones

	Protest	
	(1)	(2)
Activist Tweet $\%_{i,t-1}$	-6.28	-6.14
	(8.06)	(8.05)
Activist Retweet $\%_{i,t-1}$	-1.36	-1.91
	(6.85)	(6.90)
Activist Mention $\%_{i,t-1}$	8.40	8.02
	(8.99)	(9.00)
Activist Hashtag $\%_{i,t-1}$	-0.84	-0.50
	(1.11)	(1.10)
Activist Link $\%_{i,t-1}$	10.90	7.93
	(11.71)	(11.73)
$\text{Protest}_{i,t-1}$	0.01^{*}	0.01
	(0.01)	(0.01)
$\text{Repression}_{i,t-1}$	0.01	0.01
	(0.02)	(0.02)
Non-Activist $\text{Coordination}_{i,t-1}$	5.44^{***}	5.60^{***}
	(0.84)	(0.85)
Activist Coordination Tweet $\%_{i,t-1}$	-271.44^{**}	-437.98^{***}
	(116.12)	(168.53)
Mobile Phone $\%_{i,t-1}$	6.04^{*}	5.44
	(3.52)	(3.52)
Activist Coordination Tweet $\%_{i,t-1}$* Mobile Phone $\%_{i,t-1}$		4,767.43
		(4,194.30)
Intercept	-2.09^{***}	-2.17^{***}
	(0.49)	(0.49)
Country FE	Yes	Yes
N	168	168
Log Likelihood	-472.11	-471.26

$^{*}p < .1$; $^{**}p < .05$; $^{***}p < .01$

Conclusion

The aim of this chapter was to offer a broad overview of how social media data is currently being used in social science research, as well as a detailed account of the main strengths and weaknesses of this source of information about political and social behavior. Overall, our assessment shows the immense and still largely untapped potential of social media as data generators. This means we are still at an early stage in the development of standards and best practices in the different literatures that increasingly rely on this type of data; but also that we should expect to see much exciting new research being published over the next few years.

One particular challenge that we didn't discuss at length is the extent to which research that uses social media data can yield findings that can be generalizable across domains and over time, as opposed to idiosyncratic to the particular site whose data is being used. It is certainly the case that some of the websites we study did not even exist ten years ago, which begs the question of whether they will still exist ten years from now. We have plenty of examples of successful sites that eventually disappeared, such as MySpace or Friendster. However, our view is that even if Facebook or Twitter eventually disappear, findings derived from research on these sites will survive and remain valid. Contrary to the common view that characterizes social media interactions as not occurring in the "real world", behavior on these sites indeed mirrors offline behavior, and thus it can help reveal the mechanisms that drive human behavior, not only on these platforms, but in people's lives more generally.

References

Allcott, Hunt and Matthew Gentzkow. 2017. "Social media and fake news in the 2016 election." *Journal of Economic Perspectives* 31(2):211–36.

Alsaedi, Nasser, Pete Burnap and Omer Rana. 2017. "Can We Predict a Riot? Disruptive Event Detection Using Twitter." *ACM Transactions on Internet Technology* 17(2):1–26.

Amar, Paul. 2011. "Turning the Gendered Politics of the Security State Inside Out? Charging the Police with Sexual Harassment in Egypt." *International Feminist Journal of Politics* 13(3):299–328.

Anastasopoulos, Lefteris and Jake Williams. 2017. "Understanding collective action with Bayesian social action identification algorithms .".

Bakshy, Eytan, Solomon Messing and Lada A Adamic. 2015. "Exposure to ideologically diverse news and opinion on Facebook." *Science* 348(6239):1130–1132.

Barbera, Pablo. 2015. "Birds of the Same Feather Tweet Together: Bayesian Ideal Point Estimation Using Twitter Data." *Political Analysis* 23(August 2013):76–91.

Barberá, Pablo and Gonzalo Rivero. 2015. "Understanding the political representativeness of Twitter users." *Social Science Computer Review* 33(6):712–729.

Barberá, Pablo, John T Jost, Jonathan Nagler, Joshua A Tucker and Richard Bonneau. 2015. "Tweeting from left to right: Is online political communication more than an echo chamber?" *Psychological science* 26(10):1531–1542.

Barberá, Pablo, Ning Wang, Richard Bonneau, John T Jost, Jonathan Nagler, Joshua Tucker and Sandra González-Bailón. 2015. "The critical periphery in the growth of social protests." *PloS one* 10(11):e0143611.

Barberá, Pablo, Richard Bonneau, Patrick Egan, John T Jost, Jonathan Nagler and Joshua Tucker. 2014. Leaders or followers? Measuring political responsiveness in the US Congress using social media data. In *110th American Political Science Association Annual Meeting*.

Barberá, Pablo and Thomas Zeitzoff. 2017. "The new public address system: why do world leaders adopt social media?" *International Studies Quarterly* 62(1):121–130.

Barnidge, Matthew. 2017. "Exposure to political disagreement in social media versus face-to-face and anonymous online settings." *Political Communication* 34(2):302–321.

Beauchamp, Nicholas. 2017. "Predicting and interpolating state-level polls using Twitter textual data." *American Journal of Political Science* 61(2):490–503.

Beinin, Joel. 2009. Workers' Struggles Under 'Socialism' and Neoliberalism. In *Egypt: The Moment of Change*, ed. Philip Marfleet and Rabab El-Mahdi. London: Zed Books pp. 68–86.

Bisbee, James and Jennifer M Larson. 2017. "Testing Social Science Network Theories with Online Network Data: An Evaluation of External Validity." *American Political Science Review* 111(3):502–521.

Bliss, Catherine A., Isabel M. Kloumann, Kameron Decker Harris, Christopher M. Danforth and Peter Sheridan Dodds. 2012. "Twitter reciprocal reply networks exhibit assortativity with respect to happiness." *Journal of Computational Science* 3(5):388–397.

Bond, Robert M., Christopher J. Fariss, Jason J. Jones, Adam D.I. Kramer, Cameron Marlow, Jaime E. Settle and James H Fowler. 2012. "A 61-million-person experiment in social influence and political mobilization." *Nature* 489(7415):295–8.

Bond, Robert and Solomon Messing. 2015. "Quantifying Social Media's Political Space: Estimating Ideology from Publicly Revealed Preferences on Facebook." *American Political Science Review* 109(1):62–78.

Borge-Holthoefer, Javier, Walid Magdy, Kareem Darwish and Ingmar Weber. 2015. Content and Network Dynamics Behind Egyptian Political Polarization on Twitter. In *18th Conference on Computer-Supported Cooperative Work and Social Computing*. pp. 1–30.

Boschee, Elizabeth, Jennifer Lautenschlager, Sean O'Brien, Steve Shellman, James Starz and Michael Ward. 2015. "ICEWS Coded Event Data.".
URL: *http://dx.doi.org/10.7910/DVN/28075*

Boxell, Levi, Matthew Gentzkow and Jesse M Shapiro. 2017. "Greater Internet use is not associated with faster growth in political polarization among US demographic groups." *Proceedings of the National Academy of Sciences* p. 201706588.

Breuer, Anita. 2012. The Role of social media in mobilizing political protest: Evidence from the Tunisian Revolution. Technical report German Development Institute Bonn: .

Bright, Jonathan. 2018. "Explaining the emergence of political fragmentation on social media: the role of ideology and extremism." *Journal of Computer-Mediated Communication* 23(1):17–33.

Broockman, David E and Donald P Green. 2014. "Do online advertisements increase political candidates name recognition or favorability? Evidence from randomized field experiments." *Political Behavior* 36(2):263–289.

Brym, Robert, Melissa Godbout, Andreas Hoffbauer, Gabe Menard and Tony Huiquan Zhang. 2014. "Social media in the 2011 Egyptian uprising." *The British journal of Sociology* 65(2):266–92.

Cadena, Jose, Gizem Korkmaz, Chris J Kuhlman, Achla Marathe, Naren Ramakrishnan and Anil Vullikanti. 2015. "Forecasting Social Unrest Using Activity Cascades." *PloS One* 10(6):e0128879.

Center, Pew Research. 2017. "Partisan Conflict and Congressional Outreach." *Pew Research Center Report* .

Ceron, Andrea, Luigi Curini, Stefano M Iacus and Giuseppe Porro. 2014. "Every tweet counts? How sentiment analysis of social media can improve our knowledge of citizens political preferences with an application to Italy and France." *New Media & Society* 16(2):340–358.

Chwe, Michael. 1998. "Believe the Hype: Solving Coordination Problems with Television Advertising.".

Coakley, Kevin, Christine Kirkpatrick and Zachary C. Steinert-Threlkeld. 2019. "rehydratoR.".
URL: *https://cran.r-project.org/web/packages/rehydratoR/index.html*

Coppock, Alexander, Andrew Guess and John Ternovski. 2016. "When Treatments are Tweets: A Network Mobilization Experiment over Twitter." *Political Behavior* 38(1):105–128.

Driscoll, Jesse and Zachary C. Steinert-Threlkeld. 2018. "Does Social Media Enable Irredentist Information Warfare?".

Dunbar, R.I.M., Valerio Arnaboldi, Marco Conti and Andrea Passarella. 2015. "The structure of online social networks mirrors those in the offline world." *Social Networks* 43:39–47.

Dwork, Cynthia. 2008. Differential privacy: A survey of results. In *International Conference on Theory and Applications of Models of Computation*. Springer pp. 1–19.

Eckles, Dean, Brett R Gordon and Garrett A Johnson. 2018. "Field studies of psychologically targeted ads face threats to internal validity." *Proceedings of the National Academy of Sciences* p. 201805363.

Ferrara, Emilio, Onur Varol, Clayton Davis, Filippo Menczer and Alessandro Flammini. 2016. BotOrNot: A System to Evaluate Social Bots. In *Proceedings of the 25th International Conference Companion on World Wide Web*. pp. 273–274.

Fletcher, Richard and Rasmus Kleis Nielsen. 2018. "Are people incidentally exposed to news on social media? A comparative analysis." *New Media & Society* 20(7):2450–2468.

Freelon, D., M. Lynch and S. Aday. 2015. "Online Fragmentation in Wartime: A Longitudinal Analysis of Tweets about Syria, 2011-2013." *The ANNALS of the American Academy of Political and Social Science* 659(1):166–179.

Freelon, Deen. 2018. "Computational research in the post-API age." *Political Communication* pp. 1–4.

Freelon, Deen, Charlton McIlwain and Meredith Clark. 2016. "Beyond the hashtags:# Ferguson,# Blacklivesmatter, and the online struggle for offline justice.".

Gayo Avello, Daniel, Panagiotis T Metaxas and Eni Mustafaraj. 2011. Limits of electoral predictions using twitter. In *Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media*. Association for the Advancement of Artificial Intelligence.

Gerner, Deborah J., Philip A. Schrodt, Rajaa Abu-Jabr and Omur Yilmaz. 2002. Conflict and Mediation Event Observations (CAMEO): A New Event Data Framework for the Analysis of Foreign Policy Interactions. In *Annual Meeting of the International Studies Association*.

Ghannam, Jeffrey. 2011. Social Media in the Arab World: Leading up to the Uprisings of 2011. Technical report Center for International Media Assistance Washington D.C.: .

Gladwell, Malcolm. 2010. "Small change.".

Golder, Scott A and Michael W Macy. 2011. "Diurnal and seasonal mood vary with work, sleep, and daylength across diverse cultures." *Science (New York, N.Y.)* 333(6051):1878–81.

Golder, Scott A and Michael W Macy. 2014. "Digital footprints: Opportunities and challenges for online social research." *Annual Review of Sociology* 40:129–152.

Gosling, Samuel D, Adam a Augustine, Simine Vazire, Nicholas Holtzman and Sam Gaddis. 2011. "Manifestations of personality in Online Social Networks: self-reported Facebook-related behaviors and observable profile information." *Cyberpsychology, behavior and social networking* 14(9):483–8.

Grimmer, Justin, Solomon Messing and Sean J Westwood. 2012. "How words and money cultivate a personal vote: The effect of legislator credit claiming on constituent credit allocation." *American Political Science Review* 106(4):703–719.

Gruzd, Anatoliy and Ksenia Tsyganova. 2015. "Information Wars and Online Activism During the 2013/2014 Crisis in Ukraine: Examining the Social Structures of Pro- and Anti-Maidan Groups." *Policy & Internet* 7(2):121–158.
URL: *http://doi.wiley.com/10.1002/poi3.91*

Guess, Andrew, Jonathan Nagler and Joshua Tucker. 2019. "Less than you think: Prevalence and predictors of fake news dissemination on Facebook." *Science advances* 5(1):eaau4586.

Gunning, Jeroen and Ilan Zvi Baron. 2013. *Why Occupy a Square: People, Protests and Movements in the Egyptian Revolution*. London: Hurst Publishers.

Hobbs, William R and Margaret E Roberts. 2018. "How sudden censorship can increase access to information." *American Political Science Review* pp. 1–16.

Hsuan, Ted, Yun Chen, Paul Zachary and Christopher J Fariss. 2017. "Who Protests? Using Social Media Data to Estimate How Social Context Affects Political Behavior.".

Huckfeldt, Robert, Paul E Johnson and John Sprague. 2004. *Political disagreement: The survival of diverse opinions within communication networks*. Cambridge University Press.

Jamieson, Kathleen Hall. 2018. *Cyberwar: How Russian Hackers and Trolls Helped Elect a President What We Don't, Can't, and Do Know*. Oxford: Oxford University Press.

Jones, Jason J, Jaime E Settle, Robert M Bond, Christopher J Fariss, Cameron Marlow and James H Fowler. 2013. "Inferring tie strength from online directed behavior." *PloS one* 8(1):e52168.

Joo, Jungseock, Donghyeon Won and Zachary C. Steinert-Threlkeld. 2018. "Mechanisms of Collective Action: Measuring the Role of Violence, Identity, and Free-riding Using Geolocated Images.".

Joo, Jungseock, L. Jason Anastasopoulos and Zachary C. Steinert-Threlkeld. 2019. "Understanding the Political Ideology of Legislators from Social Media Images.".

Kalla, Joshua L and David E Broockman. 2018. "The minimal persuasive effects of campaign contact in general elections: Evidence from 49 field experiments." *American Political Science Review* 112(1):148–166.

Keller, Franziska B, David Schoch, Sebastian Stier and JungHwan Yang. 2017. How to Manipulate Social Media: Analyzing Political Astroturfing Using Ground Truth Data from South Korea. In *ICWSM*. pp. 564–567.

King, Gary, Jennifer Pan and Margaret E Roberts. 2013. "How censorship in China allows government criticism but silences collective expression." *American Political Science Review* 107(2):326–343.

King, Gary, Jennifer Pan and Margaret E Roberts. 2014. "Reverse-engineering censorship in China: Randomized experimentation and participant observation." *Science* 345(6199):1251722.

King, Gary, Jennifer Pan and Margaret E Roberts. 2017. "How the Chinese government fabricates social media posts for strategic distraction, not engaged argument." *American Political Science Review* 111(3):484–501.

Klašnja, Marko, Pablo Barberá, Nick Beauchamp, Jonathan Nagler and Joshua Tucker. 2016. Measuring public opinion with social media data. In *The Oxford Handbook of Polling and Survey Methods*.

Kosinski, Michal, David Stillwell and Thore Graepel. 2013*a*. "Private traits and attributes are predictable from digital records of human behavior." *Proceedings of the National Academy of Sciences of the United States of America* 110(15):5802–5.

Kosinski, Michal, David Stillwell and Thore Graepel. 2013*b*. "Private traits and attributes are predictable from digital records of human behavior." *Proceedings of the National Academy of Sciences* p. 201218772.

Kostyuk, Nadiya and Yuri M. Zhukov. 2017. "Invisible Digital Front: Can Cyber Attacks Shape Battlefield Events?" *Journal of Conflict Resolution* 20(10):1–31.

Kramer, Adam DI, Jamie E Guillory and Jeffrey T Hancock. 2014. "Experimental evidence of massive-scale emotional contagion through social networks." *Proceedings of the National Academy of Sciences* p. 201320040.

Kreiss, Daniel. 2012. *Taking our country back: The crafting of networked politics from Howard Dean to Barack Obama*. Oxford University Press.

Kreiss, Daniel. 2016. "Seizing the moment: The presidential campaigns use of Twitter during the 2012 electoral cycle." *New Media & Society* 18(8):1473–1490.

Langohr, Vickie. 2013. ""This is Our Square": Fighting Sexual Assault at Cairo Protests." *Middle East Research and Information Project* 43(Fall).
URL: *http://www.merip.org/mer/mer268/our-square*

Lerman, Kristina, Rumi Ghosh and Tawan Surachawala. 2012. "Social contagion: An empirical study of information spread on Digg and Twitter follower graphs." *arXiv preprint arXiv:1202.3162* .

Lim, Merlyna. 2012. "Clicks, Cabs, and Coffee Houses: Social Media and Oppositional Movements in Egypt, 2004-2011." *Journal of Communication* 62(2):231–248.

Little, Andrew T. 2015. "Communication Technology and Protest." *Journal of Politics* 78(1):152–166.

Malik, Momin M, Hemank Lamba, Constantine Nakos and Jurgen Pfeffer. 2015. Population Bias in Geotagged Tweets. In *9th International AAAI Conference on Weblogs and Social Media*. pp. 18–27.

Moreno, Jacob Levy. 1934. "Who shall survive?: A new approach to the problem of human interrelations.".

Morozov, Evgeny. 2012. *The net delusion: The dark side of Internet freedom*. PublicAffairs.

Morstatter, Fred, Jürgen Pfeffer, Huan Liu and Kathleen M Carley. 2013. Is the Sample Good Enough? Comparing Data from Twitter's Streaming API with Twitter's Firehose. In *ICWSM*.

Munger, Kevin. 2016. "Tweetment Effects on the Tweeted: Experimentally Reducing Racist Harassment." *Political Behavior* pp. 1–21.

Munger, Kevin, Richard Bonneau, Jonathan Nagler and Joshua A. Tucker. 2018. "Elites Tweet to get Feet off the Streets: Measuring Regime Response to Protest Using Social Media." *Political Science Research and Methods* First View:1–20.

Mutz, Diana C. 2002. "Cross-cutting social networks: Testing democratic theory in practice." *American Political Science Review* 96(1):111–126.

Nulty, Paul, Yannis Theocharis, Sebastian Adrian Popa, Olivier Parnet and Kenneth Benoit. 2016. "Social media and political communication in the 2014 elections to the European Parliament." *Electoral studies* 44:429–444.

O'Callaghan, Derek, Nico Prucha, Derek Greene, Maura Conway, Joe Carthy and Padraig Cunningham. 2014. "Online Social Media in the Syria Conflict: Encompassing the Extremes and the In-Betweens.".

Pariser, Eli. 2011. *The filter bubble: What the Internet is hiding from you*. Penguin UK.

Pfeffer, Jürgen and Katja Mayer. 2018. "Tampering with Twitter's Sample API." *EPJ Data Science* 7(50):1–21.

Putnam, Robert D. 2000. Bowling alone: America's declining social capital. In *Culture and politics*. Springer pp. 223–234.

Raab, Gillian M, Beata Nowok and Chris Dibben. 2016. "Practical data synthesis for large samples." *Journal of Privacy and Confidentiality* 7(3):67–97.

Radford, Jason and Betsy Sinclair. 2016. "Electronic Homestyle: Tweeting Ideology.".

Sakaki, Takeshi, Makoto Okazaki and Yutaka Matsuo. 2010. Earthquake Shakes Twitter Users: Real-time Event Detection by Social Sensors. In *International World Wide Web Conference*. pp. 851–860.

Schrodt, Philip A., Shannon G. Davis and Judith L. Weddle. 1994. "Political Science: KEDS–A Program for the Machine Coding of Event Data." *Social Science Computer Review* 12(4):561–587.

Siegel, Alexandra A. and Vivienne Badaan. 2018. "#No2Sectarianism: Experimental Approaches to Reducing Sectarian Hate Speech Online.".

Steinert-Threlkeld, Zachary C. 2017*a*. "Longitudinal Network Analysis with Incomplete Data." *Political Analysis* .

Steinert-Threlkeld, Zachary C. 2017*b*. "Spontaneous Collective Action: Peripheral Mobilization During the Arab Spring." *American Political Science Review* 111(02):379–403.

Steinert-Threlkeld, Zachary C. 2018. *Twitter as Data*. Cambridge University Press.

Steinert-Threlkeld, Zachary C. 2019. "Comment: The Future of Event Data is Images." *Sociological Methodology* 49.

Steinert-Threlkeld, Zachary C. and Jungseock Joo. 2019. "Event Data with Images.".

Stier, Sebastian, Lisa Posch, Arnim Bleier and Markus Strohmaier. 2017. "When populists become popular: comparing Facebook use by the right-wing movement Pegida and German political parties." *Information, Communication & Society* 20(9):1365–1388.

Sunstein, Cass R. 2018. *# Republic: Divided democracy in the age of social media*. Princeton University Press.

Takhteyev, Yuri, Anatoliy Gruzd and Barry Wellman. 2012. "Geography of Twitter networks." *Social Networks* 34(1):73–81.

Theocharis, Yannis, Pablo Barberá, Zoltán Fazekas, Sebastian Adrian Popa and Olivier Parnet. 2016. "A bad workman blames his tweets: the consequences of citizens' uncivil Twitter use when interacting with party candidates." *Journal of communication* 66(6):1007–1031.

Timoneda, Joan. 2018. "Where in the world is my tweet: Detecting irregular removal patterns on Twitter." *PloS one* 13(9):e0203104.

Tufekci, Zeynep and Christopher Wilson. 2012. "Social Media and the Decision to Participate in Political Protest: Observations From Tahrir Square." *Journal of Communication* 62(2):363–379.

Tumasjan, Andranik, Timm Oliver Sprenger, Philipp G Sandner and Isabell M Welpe. 2010. "Predicting elections with twitter: What 140 characters reveal about political sentiment." *Icwsm* 10(1):178–185.

Ugander, Johan, Brian Karrer, Lars Backstrom and Cameron Marlow. 2011. "The anatomy of the facebook social graph." *arXiv preprint arXiv:1111.4503* .

Valenzuela, Sebastián. 2013. "Unpacking the use of social media for protest behavior: The roles of information, opinion expression, and activism." *American Behavioral Scientist* 57(7):920–942.

Wilson, R. E., S. D. Gosling and L. T. Graham. 2012. "A Review of Facebook Research in the Social Sciences." *Perspectives on Psychological Science* 7(3):203–220.

Wilson, Steven Lloyd. 2017. "Detecting Mass Protest through Social Media." *Journal of Social Media in Society* 6(2):5–25.

Youyou, Wu, Michal Kosinski and David Stillwell. 2015. "Computer-based personality judgments are more accurate than those made by humans." *Proceedings of the National Academy of Sciences* 112(4):1036–1040.

Zachary, Wayne W. 1977. "An information flow model for conflict and fission in small groups." *Journal of anthropological research* 33(4):452–473.

Zamal, Faiyaz Al, Wendy Liu and Derek Ruths. 2012. Homophily and Latent Attribute Inference: Inferring Latent Attributes of Twitter Users from Neighbors. In *Sixth International AAAI Conference on Weblogs and Social Media*. pp. 387–390.

Zeitzoff, Thomas. 2011. "Using Social Media to Measure Conflict Dynamics: An Application to the 2008-2009 Gaza Conflict." *Journal of Conflict Resolution* 55(6):938–969.

Zeitzoff, Thomas. 2016. "Does Social Media Influence Conflict? Evidence from the 2012 Gaza Conflict." *Journal of Conflict Resolution* pp. 1–35.

Zeitzoff, Thomas, John Kelly and Gilad Lotan. 2015. "Using social media to measure foreign policy dynamics: An empirical analysis of the Iranian-Israeli confrontation (2012-13)." *Journal of Peace Research* 52(3):368–383.

Zhang, Han and Jennifer Pan. 2019. "CASM: A Deep-Learning Approach for Identifying Collective Action Events with Text and Image Data from Social Media." *Sociological Methodology* 49:1–48.

Zhukov, Yuri M. 2015. "Trading hard hats for combat helmets: The economics of rebellion in eastern Ukraine." *Journal of Comparative Economics* 44(1):1–15.

Zimmer, Michael. 2010. "But the data is already public: on the ethics of research in Facebook." *Ethics and information technology* 12(4):313–325.