Future protest made risky: Examining social media based civil unrest prediction research and products

Gabriel Grill[1]

[1]*University of Michigan, Ann Arbor, MI, USA (E-mail: ggrill@umich.edu)*

Abstract. Social media has both been hailed for enabling social movements and critiqued for its affordances as a surveillance infrastructure. In this work, I focus on the latter by analyzing research, products, and discourses around the recent history of civil unrest prediction based on social media data and other public data sources, thereby giving insights into current and often opaque protest surveillance and forecasting practices. Technologies to monitor individuals and groups online have been developed for instance to predict US protests following the election of President Trump in 2016 and labor strikes across global supply chains. These works are part of an emerging computer science research field focused on "civil unrest prediction" dedicated to forecasting protests across the globe (e.g., Indonesia, Brazil, and Australia). Foremost I focus on scholarly literature as my unit of analysis, but also other artifacts discussing or detailing applications for companies, organizations or governments are examined. I provide a conceptualization of civil unrest prediction technology by illustrating data sources, features and methods used, and how prediction and detection are necessarily entangled. Then I show how various kinds of unrest activity are framed as risks to be fixed or averted for various actors with differing interests such as the military, law enforcement, and various industries. Finally, I critically unpack justifications and ascribed benefits of the technology and point to how the perspectives of protestors are almost completely absent. My analysis shows a critical need for regulation centering activists and workers, and reflection within academia, particularly in the fields of computer and data science, on the ethics and politics of protest research and ensuing technological applications.

Keywords: civil unrest prediction, labor strike, preemption, protest, risk assessment, social movements, surveillance

1. Introduction

Social media platforms have been lauded for enabling and organizing social movements, both online and offline. An extensive set of scientific literature also supports this claim (Hermida and Hernández-Santaolalla 2018; Tufekci and Wilson 2012; Valenzuela et al. 2012). It attests to how social media has been strategically used to communicate with "networked publics" (boyd 2010), i.e., groups of social media users part of online discourse, to aid in the formation of counter publics (Warner 2002). These publics often collectively voice grievances and potentially also protest offline on the ground. Such tactics are not only employed in service of emancipatory justice and the dismantling of the matrix of oppression (Costanza-Chock 2020), but also to organize and channel regressive movements such as the alt-right (Zhou et al. 2018) and the manosphere (Marwick and Caplan 2018). The enticing and emancipatory imaginaries of empowerment on social media keep users on these platforms. Still, the userbase also makes platforms interesting targets for algorithmic surveillance and profiling (Zuboff 2019) which may undermine this hopeful potential. Also, the pressing problems of misinformation and hate group formation on social media motivate the development of new data science based surveillance techniques (Tufekci 2017) for control. These dataveillance practices require critical scholarly attention so that the emancipatory potential of social media is not undermined.

One central target for predictive technologies that surveil and profile social movements online is labor organizing activity. For instance, such technology has been used to anticipate labor unrest in Walmart stores (Peterson 2015). Another example is Wholefoods (Peterson 2020) reportedly using heat maps to visualize unionization risk scores based on variables such as "racial diversity, employee loyalty, 'tipline' calls, and violations recorded by the Occupational Safety and Health Administration." According to insiders, "tracking active or potential unionization is a common practice among large companies" (Peterson 2020). This frequent usage illustrates the importance of more worker-centered oversight and research in this space. Another central area for this technology is national security. For instance, the US Department of Homeland Security has announced, in the aftermath of the January 6th attack on the US capitol, plans to build a warning system that monitors social media to anticipate security threats from unrest activity (Dilanian 2021). Previously, government-funded research has been conducted, for instance, to predict US protests following the election of President Trump in 2016 (Renaud et al. 2019). This work is also part of an emerging computer science research field focused on "civil unrest prediction" dedicated to forecasting protests across the globe (e.g., Indonesia, Brazil, and Australia), usually based on various public data sources. Researchers in this area often draw upon established data science and machine learning techniques such as event detection and prediction. Besides furthering academic knowledge on civil unrest and protests, the works in this field envision supporting various actors with different interests such as governments, law enforcement, companies, and human rights NGOs.

In this paper, I analyze research and discourses around the recent history of civil unrest prediction on social media platforms and thereby unpack motivations, aims, and the framing of civil unrest. Such a focus on scholarly works can provide insights into assumptions baked into early socio-technical systems before they later stabilize (Kline and Pinch 1996; Pinch and Bijker 1987) and are adopted as black-boxed and packed products or frameworks. Furthermore, unrest surveillance technologies are often not accessible for public scrutiny since they are shielded as company or governmental secretes (Pasquale 2015). In turn, studying publicly available scholarly works can provide interesting insights into a space that is otherwise difficult to access. My research questions are: How can civil unrest prediction be conceptualized? How is this research justified and motivated? I argue that both prediction and detection of civil unrest are temporally entangled and can be understood as risk assessment practices (Luhmann 2005). I highlight how they have emerged over the last years as part of the rise of techno-security culture and techno-optimistic promises (Avle et al. 2020) of big data. Finally, I discuss justifications for this transformation and argue for further research and a debate on the ethics and politics of civil unrest prediction and detection. Foremost, I focus on scholarly literature as my unit of analysis and also examine other documents discussing or detailing applications for companies, organizations, or governments.

My research into unrest surveillance and profiling, here in the form of prediction and detection, matters as it provides insights into changing risk assessment and mediation practices which reconfigure power relationships between activists, publics, states, industries, and human-rights organizations. This work employs the critical lenses of science and technology studies (STS) in its analysis. It aims to contribute to scholarly discourses in CSCW around privacy and ethics of big data and social movement research. I hope to contribute to a body of research on the uses and development of big data technologies by state bodies (Dencik et al. 2018) and companies (Uldam 2018), as well as the involvement of academia in this endeavor. I highlight that the increasing development and adoption of such unrest risk assessment technologies pose challenges to democratic participation, labor rights, and citizenship. I aim to "study up" (Marcus and Fischer 2014) and illustrate expressly also how activists are framed, which are often underrepresented in the discourses I examined. Furthermore, I hope that this work will contribute to a needed public debate.

The sheer number of journalistic articles discussing concerns about big data technologies for protest surveillance (Ahmed 2018) and labor union avoidance (Peterson 2020) attests to this.

2. Related Work

The rise of techno-security culture (Weber and Kämpf 2020) and risk assessment practices (Beck 1992; Williams 2008) over the last decades has to led to a variety of scholarly literature critiquing predictive data technologies and preemptive practices supported by them. In particular, scholarly attention has been paid to governments and law enforcement agencies as they increasingly have adopted predictive technologies for the governance of national security risks. This popularity has been attributed to prevalent securitization technoimaginaries (Weber and Kämpf 2020) and a broader shift towards proactive and preemptive security practices (Dencik et al. 2018; Hälterlein and Ostermeier 2018; Vogel et al. 2016). In particular, predictive policing techniques have surged in adoption and, in consequence, were critiqued by various scholars for reifying inequalities (Ferguson 2017; Richardson et al. 2019). They have been employed to supposedly improve efficiency/effectivity and preemptively combat criminal activity while framed to varying degrees as success stories in the short-term through reports and experiments. These systems calculate the riskiness of criminality over time for locations, individuals, and groups by consolidating various databases, often on previous arrests, to mine them for patterns. The consequences of being marked as risky vary. Places may be policed more frequently, while individuals may become subject to heightened surveillance. These approaches have been critiqued for reproducing and amplifying systemic racism and other forms of inequality prevalent in contemporary policing practices (Richardson et al. 2019) and broader society. Scholars have argued that these systems support the production of feedback loops (Ensign et al. 2017; Gandy 2016), leading to evermore policing of marginalized populations, particularly people of color and the poor.

Historically collected social data part of governmental databases such as family ties have been used to calculate risk profiles (Ferguson 2017), but also online social media activity data have recently been harvested and analyzed for this purpose (Gerber 2014; Williams et al. 2017). An examination (Pelzer 2018) of research on the prediction of inclination towards terrorist activities or radicalization on social media and, in particular, also its limitations, concluded that it is currently not clear in what capacity police use machine learning tools. The authors argued against their applicability in practice due to their narrow focus on the accuracy of pattern recognition. A survey conducted in 2016 found that social media surveillance tools were employed in investigations by 89% of US police departments (Borradaile et al. 2020), highlighting their widespread use. The technical capabilities of these tools, though remain mostly unclear. A study (Borradaile et al. 2020) of a reverse engineered social media surveillance tool from the Corvallis (Oregon) Police Department found that it employed a simple keyword search system, which seemingly was based on terms almost not related to the topic of interest to the officers. The authors concluded that this calls "into question the utility that such a keyword based search could have to law enforcement" (p. 1). Social media surveillance tools for policing developed by companies, such as MediaSonar and Geofeedia, were reportedly used also to arrest and track #BlackLivesMatter protestors. However, after public critique from the civil-rights organization ACLU (American Civil Liberties Union), both companies lost their social media API access rights (Borradaile et al. 2020).

An interview-based study (Dencik et al. 2018) on the use of social media data in protest policing in the UK conducted between August and September 2015 found that law enforcement agencies predominantly employed commercial marketing tools. No technologies developed explicitly for online protest surveillance were used. The authors raise concerns on the current setup due to various

forms of bias creep, the inherently limited certainty of protest anticipation based on social media data, and the lacking transparency of commercial tools. An analysis (Egawhary 2019) of internal policy documents on the use of social media acquired through the freedom of information act in the UK has concluded that the police mainly use social media analysis tools for PR and online advertising. This finding aligns with other work (Colbran 2018) that highlights how social media has become a tool for police departments to manage their image and frame their work to the public in positive ways, in turn circumventing the press. These developments show an increasing need to pay attention to tools for social media surveillance regarding policing and message control. Information campaigning online to combat protest formation early on and divert attention has also become an essential tool to repressive governments (Tufekci 2017).

In contrast to civil unrest surveillance practices of governments and law enforcement agencies on social media, companies remain understudied (Uldam 2018). The author highlights that reputational risks matter significantly for companies, driving the hiring of risk assessment and PR agencies. Consequently, surveillance focuses on groups that aim to use their voice for critique, thereby holding companies accountable. The information collected online is employed to create strategies to manage criticism and remove it from the public's eye. Other prior work on predictive technologies targeting civil unrest has also argued for a need for further research in this area (Grill 2020) and highlighted possible methodological approaches for such an endeavor (Heimstädt and Dobusch 2021). This literature review highlights that a significant amount of critical scholarly work on predictive policing has been conducted, and it has illustrated various social and technical problems with this technology. However, few works focus specifically on civil unrest prediction. As a result, the practices and products of predictive surveillance companies targeting social media have received little attention. I contribute to these scholarly debates on affordances of and discourses around surveillance technologies by examining civil unrest prediction research, particularly of protests and (labor) strikes on social media.

2. Methods

The widespread usage of social media platforms and the increasing availability of online public data has motivated the creation of research and products concerned with the prediction and risk assessment of civil unrest. I employ situational analysis (Clarke et al. 2017) to make sense of this recent turn and study this situation by considering the various actors and elements that include research publications, products, positions, techno-optimistic promises, discourses, academic institutions, companies, and states. This method provides three cartographic approaches to map out complex situations while encouraging sensibilities to include different standpoints, especially ones marginalized or seemingly absent. Moreover, situational analysis is compatible with an interpretivist analytical framework grounded in science and technology studies (STS) and the co-production of science, technology, and social order (Jasanoff 2004).

2.1. Data Collection

Foremost, my analysis is grounded in the recent scholarly computer science literature on civil unrest prediction due to its recent concern with unrest surveillance technologies and public availability. In contrast, only little documentation is available on tools directly used by law enforcement agencies and industry. The research paper collection was conducted in August 2019. I have crafted several queries to search through multiple academic literature databases and retrieved an initial set of publications on civil unrest prediction. The terms I used for this inquiry were "protest," "social unrest," "labor strike," and "civil unrest" in conjunction with "social media." My search encompassed

all publications whose title, keywords, or abstract match the search strings. I did not use "prediction" as an additional qualifying term in this initial search because I wanted the query to yield more results, and I was unsure if all researchers would use the term prediction, as terms such as forecast and detection were sometimes used synonymously.

I retrieved publications in my initial search from the IEEE and ACM databases, which are major literature databases in computer science. From this set, I classified a total of 33 papers as concerned with variants of offline civil unrest prediction based on public data sources. I added 20 additional publications from other academic venues to my corpus by examining works citing the initial set of publications on Google Scholar. My intention for including these additional publications was not to provide a complete overview of all research on civil unrest prediction but rather to better understand the scholarly discourse in which these initial articles from major computer science venues are embedded. The publications all describe systems for anticipating future unrest in some capacity through online social data, which in my definition encompasses social media, blogs, and also Google Trends. As further elaborated in section 3.1., the distinction between the detection of ongoing unrest in the present and prediction of future unrest is not always clear. In turn, I classified several papers on unrest detection as also concerned with predicting or anticipating unrest and included them in the corpus. The search also yielded several publications that did not fit my topic of interest, such as works concerned specifically with crime prediction or ones that only considered news sites as data sources. In total, I acquired 53 publications in English for my principal analysis. In table 1, I list the publications I have identified and their publication date.

Year	Publications	#
2010	Colbaugh and Glass 2010	1
2013	Compton et al. 2013; Hua, Chen, et al. 2013; Hua, Lu, et al. 2013; Manrique et al. 2013; Tan et al. 2013; Williams et al. 2013	6
2014	Benkhelifa et al. 2014; Chen and Neill 2014; Compton et al. 2014; Doyle et al. 2014; Filchenkov et al. 2014; Kallus 2014; Ramakrishnan et al. 2014; Xu et al. 2014	8
2015	Agarwal et al. 2015; Agarwal and Sureka 2015; Cadena et al. 2015; Chen and Neill 2015; Goode et al. 2015; Hoegh et al. 2015; Korkmaz et al. 2015; Muthiah et al. 2015; Ramakrishnan et al. 2015; Zhao et al. 2015	10
2016	Agarwal and Sureka 2016; De Choudhury et al. 2016; Ganar and Ardhapurkar 2016; Korkmaz et al. 2016; Korolov et al. 2016; Liu et al. 2016; Muthiah, Huang, et al. 2016; Muthiah, Vullikanti, et al. 2016; Qi et al. 2016; Zhao et al. 2016	10
2017	Agarwal 2017; Alsaedi et al. 2017; Gupta et al. 2017; Kang et al. 2017; Mishler et al. 2017; Osborne et al. 2017; Wang et al. 2017; Wilson 2017; Zhao et al. 2017	9
2018	Ansah et al. 2018a, 2018b; Bahrami et al. 2018; Hossny and Mitchell 2018; Rule et al. 2018; Singh and Pal 2018; Wu and Gerber 2017	7
2019	Ertugrul et al. 2019; Renaud et al. 2019	2

Table 1. Publications examined in the situational analysis.

In order to get a broader perspective on the situation of civil unrest prediction research, I also collected other relevant artifacts in English through an explorative internet search. The artifact collection encompasses reports, articles, webpages, videos, and presentations pertinent to civil unrest prediction products or practices of companies, research institutes and, (non-)governmental organizations. These artifacts are not necessarily intended for an academic audience and thereby differ in quality but still provide interesting partial perspectives that support me in answering my research questions. I conducted the data collection by submitting crafted queries based on keywords used for my previous search of academic literature and ones tailored explicitly to application areas

such as risk assessment in supply chain management to a popular search engine. I found several articles, videos, and websites on startups, companies and, other organizations concerned with civil unrest prediction through social media data. In total, I considered 51 artifacts collected between August 2019 and May 2020.

2.2. Analysis

My situational analysis is based on a qualitative multi-method approach. In the first phase, I conducted a Document Analysis (DA) (Bowen 2009), which is an iterative process consisting of "finding, selecting, appraising (making sense of), and synthesizing data contained in documents" (Bowen 2009). It involves a coding step, directed qualitative content analysis (Hsieh and Shannon 2005), and thematic analysis (Williamson et al. 2018). Following situational analysis, it also includes an interpretation of purpose, context, completeness, and target audience of the documents (Bowen 2009). My research considers studied documents as performative and reductive descriptions of complex socio-technical artifacts developed over long periods of time which do not simply mirror underlying realities. A similar approach based on thematic analysis of academic literature has been employed to study violence of misgendering automatic gender recognition systems (Keyes 2018).

The results of the DA function as a basis for my second phase, which encompasses mapping the situation and critical discourse analysis (CDA) (Gee 2014; Mullet 2018). I have used the extracted codes and themes of the documents to construct various maps to visualize positions, discourses and power relations, and implicated actors. The collected artifacts were used in this mapping process to better contextualize the situation, for instance, by including spin-off companies that arose out of certain research projects and highlighting funding relationships. CDA is considered a part of the situation analysis framework as the maps also aid in understanding discourses through visualization. A stepwise approach involving coding, theme extraction, and CDA has been employed, e.g., to study statements of Mark Zuckerberg on Facebook (Hoffmann et al. 2018). Also, the General Analytic Framework for CDA (Mullet 2018) has been exemplified on academic literature (p. 125), which further illustrates its utility for the study of academic works, such as civil unrest prediction. By using sensibilities from CDA and situational analysis, I acknowledge that the study of discourses matters as they "construct, maintain, and legitimize social inequalities" (Mullet 2018) and techno-politics. Discourses are acts "always part and parcel of, and partially constitutive of, specific social practices" (Gee 2014). Situational analysis and CDA aid me in unpacking debates and power relations that situate, form, and stabilize civil unrest prediction practices in society. In the following sections, I first characterize civil unrest prediction, then highlight how it can be understood as a risk assessment practice, and finally unpack presented justifications for this research.

3. Conceptualization of Civil Unrest Prediction

This section describes civil unrest prediction technologies and research based on my data and illustrates some inherent difficulties and ambiguities in this endeavor. The majority of the analyzed works have been conducted by researchers affiliated with US institutions or companies, which is to be expected also because the ACM and IEEE research databases were used as a starting point, and they feature a significant amount of US research. Other more frequent affiliations include the UK, Australia, and India. The overwhelming majority of the research I considered originates in the Global North and targets national and international unrest worldwide. There are also certain areas, such as Latin America, which have received considerable attention in the literature due to the focus of big projects like the US government funded EMBERS (Muthiah, Vullikanti, et al. 2016) project. In future work, I aim to map out and discuss in more detail what places and groups are particularly targeted

by specific research efforts. In computer and data science, civil unrest prediction is often considered to be a form of analysis and anticipation of "social events" (Zhao et al. 2016, p. 3), usually based on "big public data" (Kallus 2014; Xu et al. 2014) such as social media activity or economic indicators (Korkmaz et al. 2016, p. 2). Although most of the research and products I have investigated are concerned with events, some in my corpus also focus on predicting ascribed characteristics of social movements such as their endurance, sustainability (Colbaugh and Glass 2010), and vitality (Tan et al. 2013) over time. There are two main approaches to track such events across time: detection and forecasting (Zhao et al. 2016, p. 3). I am mainly concerned with the latter, the anticipation and prediction of future unrest. However, they do not make up a strict dichotomy, and in turn, both are relevant and came up frequently in the data. This section unpacks this dichotomy and highlights how contemporary civil unrest prediction systems are part of big data regimes.

3.1. The temporal entanglement of detection and prediction

The difference between detection and prediction can arguably be best understood temporally. Detection usually is concerned with the present, for instance, recognizing an ongoing event. In contrast, prediction produces claims about the future, for example, forecasting when a future event occurs. This distinction came up in many of the papers I analyzed, but it is important to note that in machine learning, prediction is also used as a shorthand for the inference of information from data, which also encompasses various so-called detection systems. In this work, prediction refers specifically to forecasting practices to anticipate the future, which was also commonly the case in the papers I analyzed. I found researchers often frame detection as the targeting of "ongoing" (Zhao et al. 2016, p. 3) or present unrest by "promptly discover[ing] new events as they occur" (Zhao et al. 2016, p. 3). It often comes with a promise of speed in the form of servicing information on civil unrest in 'real-time' (Wang et al. 2017, p. 1). In addition, prediction also promises the forecasting of "future unrest" (Qi et al. 2016, p. 5).

Nevertheless, the distinction between unrest in the making and occurring activity is blurry and, in turn, the difference between detection and prediction as well. In practice, assumptions about how the start and end of civil unrest are conceptualized determine where the difference lies. The start of unrest activity could be defined as when the first public announcements of a protest are made, at the point a certain number of activists are on the streets, or when media reports give credibility and importance to an event. The features chosen to mark unrest as part of the present determine the timespan a detection system targets. In many cases, this means detection is also concerned with the future as it targets an extended present. Similarly, when anticipated unrest events occur close to the present, the task of prediction becomes more akin to detection. These indeterminacies around the start and end of unrest activity also illustrate how design decisions determine the temporal bounding of civil unrest. It further highlights how research into unrest prediction also needs to consider detection. Also, technical affordances such as the choice of public data sources further influence how and when unrest is perceived. The affordances and thresholds that determine recognizability of the start and end of various unrest stages matter as they also construct what protests are noticed and, in turn, receive possibly increased attention and intervention.

I encountered in my analysis several instances where authors ascribed to their systems both detection and prediction capabilities. This can be attributed in some cases to ambiguous semantics surrounding both concepts. For example, one work was concerned with "detecting future social unrest" (Compton et al. 2013, p. 1). Another example is a paper titled "Civil Unrest Prediction: A Tumblr-Based Exploration" (Xu et al. 2014) which framed its introduced technology as an "early detection system" (p. 403) that extracts information from "relevant posts" (p. 403) such as

announcement dates for "detecting emerging civil unrest events" (p. 403). Several presented systems also incorporated both detection and prediction capabilities to target both present and future unrest. Such tracking functionality across time was, for instance, requested in the Open Source Indicators (OSI) program of IARPA, which funded research into "methods for continuous, automated analysis of publicly available data in order to anticipate and/or detect significant societal events" (IARPA 2011). This further highlights an entanglement of detection and prediction practices. The former may usually aid mainly in a reactive mode of governance. The latter in a proactive/preemptive mode by pointing at a future calculated as likely. However, both matter to actors interested in information on unrest. The trend in contemporary "techno-security culture" (Weber and Kämpf 2020) towards preemption and proactive governance also requires a reactive mode to stay in place. Both complement each other and may work/fail in certain situations, e.g., "spontaneous protests" (Filchenkov et al. 2014, p. 159) not discussed online before they occur pose a critical challenge to forecasting efforts. Furthermore, predictions become less accurate the farther away the targeted future is (Kallus 2014, p. 629), which further highlights the importance of surveillance of the present and short-term through detection. In turn, in this work I also considered certain systems where the boundary was unclear. This relation also highlights more generally that social research into predictive technologies needs to pay attention also to detection systems that may be already established or co-emerge with predictive approaches.

3.2. Adherence to Big Data paradigm

In the examined literature, many authors invoke "civil unrest" as a prediction or detection target without often explicitly defining what forms of activities it encompasses or simply by providing broad subcategories such as labor strikes or occupations (Chen and Neill 2014; Korkmaz et al. 2016; Muthiah et al. 2015). Many systems are presented as targeting various activities ranging from "small, nonviolent protests that address specific issues to events that turn into large-scale riots" (Korkmaz et al. 2016, p. 1). The target to be surveilled is often constructed as seemingly all forms of unrest capturable and predictable via public data, as also all of them could be or become risky. This broad conception aligns with the "big data" (Kitchin 2014) paradigm, which for this technology entails capturing as much online activity as possible since it could reveal unrest. Some publications also focus on certain events instead of limiting the study to specific locations and timeframes. It is important to note that particularly the selection of data sources, methods, optimization goals and parameters, and ground truth event data used to confirm results are central in shaping what unrest is targeted. This subsection provides a short overview of data sources, features, and methods to illustrate how they enact the big data paradigm.

The types of data used for the prediction systems include textual data from newspapers (Korkmaz et al. 2016; Ramakrishnan et al. 2014) and social media such as tweets (Agarwal 2017; Chen and Neill 2014; Wang et al. 2017; Zhao et al. 2016), Facebook pages (Ramakrishnan et al. 2014), blogs (Doyle et al. 2014; Korkmaz et al. 2016), Tumblr posts (Xu et al. 2014). Social media is often mined for the textual content of postings, user accounts, referenced URLs, and social network metadata such as retweets, geographic information, and follower networks to model, e.g., social ties. Although news and social media data were at the center of most systems, specifically Twitter was the most common data source, I also encountered many other public data sources. These include, for instance, reports from political event databases such as ICEWS and GDELT (Korkmaz et al. 2016), NASA satellite meteorological data (Ramakrishnan et al. 2014), statistics on the anonymous internet communication network Tor (Korkmaz et al. 2016), OpenTable reservation cancellations (Doyle et al. 2014), humidity measurements (Doyle et al. 2014), Google Flu Trends (Doyle et al. 2014), Klout

scores (Chen and Neill 2014), as well as economic and financial indicators such as exchange rates (Korkmaz et al. 2016).

The goal of combining a great variety of data sources to improve predictions highlights how popular assumptions of the "big data" paradigm (Kitchin 2014) are at the heart of the project of civil unrest prediction. The three central characteristics of big data are volume, variety, and velocity (Kitchin 2014, p. 3). They can be recognized within such systems: First, the great "size and complexity of social media" (Chen and Neill 2014, p. 1) and other incorporated data sources illustrate the high "volume" (Kitchin 2014, p. 3) characteristic of big data. Second, the heterogeneity of data sources, which seemingly extends to all kinds of data somehow publicly accessible, as illustrated in the previous paragraph, attests to the "variety" of big data (Kitchin 2014, p. 3). Third, many of the data sources also have high "velocity" (Kitchin 2014, p. 3) as, e.g., social media continuously provides new information. All of these characteristics are tied together through the promise that evermore data will yield better results (Van Dijck 2014) and "reveal a hidden mathematical order in the world" (McQuillan 2018). One author even argued that the "possibilities are infinite" (Kallus 2014, p. 625) when such large datasets are processed.

I encountered various rule and machine learning based approaches in my data, be it unsupervised, semi-supervised, or supervised, combined with frameworks and processing pipelines. The methods employed range from logistic regression (Doyle et al. 2014), keyword/hashtag counting (Xu et al. 2014), keyword dictionaries as classifiers (De Choudhury et al. 2016), and rule-based approaches (Singh and Pal 2018) to deep learning (Ertugrul et al. 2019), propagation tree analysis (Ansah et al. 2018a), random forests (Kallus 2014), topic clustering (Korolov et al. 2016) and SVMs (Korolov et al. 2016). These methods are employed to capture features of unrest on social media such as sentiment and affect in language use (Chen and Neill 2014), spikes in communication activity (Korkmaz et al. 2016, p. 50), social media engagement (Ertugrul et al. 2019, p. 9), information propagation characteristics (De Choudhury et al. 2016, p. 95), "intent to protest"(Qi et al. 2016, p. 5), mobilization (Korolov et al. 2016), psycho-linguistic distancing (De Choudhury et al. 2016), temporality (Muthiah et al. 2015) and locations (Ertugrul et al. 2019). Consequently, resulting models assume that these features are stable and can be captured at scale. This great number of features to be potentially extracted from big social data further points to the underlying adherence to big data and its assumptions in civil unrest prediction research.

This section concerned with conceptualizations of civil unrest prediction has first illustrated the temporal entanglement of prediction and detection and how both need to be considered when analyzing this broader field of unrest forecasting based on public data. Second, I have characterized methods and data in civil unrest prediction research and illustrated how contemporary civil unrest prediction adheres to the big data paradigm. Such detailed descriptions of the methods and data used are usually not publicly available for industrial and governmental applications, making their analysis often challenging to almost impossible. The exact training, testing, and ground truth data are usually not released by the researchers. In turn, analysis of the quality of the results remains difficult. Previous work has highlighted inherent limitations of similar big data applications. The validity of social media as a data source for predicting complex social phenomena has been called into question. Researchers have, for instance, pointed to how social media is not representative of the offline world (boyd and Crawford 2012; Hargittai 2015; Tufekci 2014). Also, failures of promising predictive big data systems such as Google Flu Trends (Lazer et al. 2014) have further shown the volatility of such data. Related work concerned with predicting armed conflicts has illustrated the difficulty of forecasting complex social events and warned of overpromising big data applications (Cederman and Weidmann 2017). In future work, I aim to discuss embedded political assumptions and technical

limitations and affordances of civil unrest prediction in more detail. The analyzed publications give some insights into the design of such systems and potentially into applications outside academia since researchers commonly collaborate with all kinds of actors and start spin-off companies in this space. The research promises to experimentally pioneer the technologies, which are then adopted. In the next section, I will highlight some of these connections of academia with other actors.

4. Imaginaries of civil unrest risk assessment

I highlight in this section that both the detection and prediction of civil unrest can be understood as risk assessment practices, where civil unrest is framed as a source of risk to certain actors to be made calculable (Luhmann 2005). However, in most of the analyzed literature, risk as an explicit framing was not present. The exception was research projects made explicitly for industry use, e.g., to assess the risk of labor strike disruptions in supply chains (Su and Chen 2018). One research team also argued that their system could also produce "risk ratings" (Kallus 2014, p. 630) besides predicting events. This omission of the risk frame by computer science researchers could be due to the connection of the field to event detection/prediction, which is not solely concerned with events framed as risky. The prediction frame signals a connection to this community while also possibly depoliticizing the research as some may disagree that all or some forms of unrest, including peaceful protest, should be surveilled and are risky. This section shows that stated motivations and aims for this research are often centered on risks to various powerful actors and highlights how these prediction technologies embody power relations as they frame and target specific online activities.

4.1. Civil unrest as national security risk

The researchers argued that civil unrest is a potential source of "violence and insecurity" (Qi et al. 2016, p. 2), "instability" (Korkmaz et al. 2016; Ramakrishnan et al. 2014), and also possibly a threat to "supply chain operations" (Kallus 2014, p. 627), in particular when it becomes "larger and more dangerous" (p. 627). In this sense, civil unrest prediction was framed as a 'technological fix' to these risks, a tool that "can greatly benefit [..] society such that the general public can be alerted in advance to avoid potential dangers" (Kang et al. 2017, p. 1). It was further described as a tool that aids in the identification of "threats" and to support "decision making for national security, law enforcement, and intelligence missions"(Doyle et al. 2014, p. 1) so "proactive actions to alleviate tensions and minimize disruption" (Kang et al. 2017, p. 1) can be taken. Ultimately, it promises to produce "recommendation[s] for anticipatory governance to take appropriate action before event[s]" (Singh and Pal 2018, p. 513) take an "outrageous or social disruption form" (p. 513). These quotes illustrate how it is imagined as a tool to control unrest and keep it within certain bounds of an imagined civility that was not clearly articulated in the papers. In turn, the thresholds (Amoore 2020) that determine when unrest is categorized as risky, violent, or "outrageous" (Singh and Pal 2018, p. 513) are very important for understanding the scope of this technology and when officials might intervene. However, this was not much discussed in the papers. The researchers mostly aimed to simply predict unrest, which often seemingly frames all unrest activity as equally risky and thereby prone to become dangerous. The explicit classification of national security riskiness is seemingly outsourced to officials who may determine which unrest activity requires intervention depending on the information available to them. In some cases, unrest prediction was even argued to "help the investigators/police to [..] to completely stop such activity" (Ganar and Ardhapurkar 2016, p. 1). This highlights how civil unrest prediction is imagined as a technology for preemptive security risk control and avoidance.

The described national security focus is also reflected in the funders of this research, which include a variety of security agencies such as the Air Force (Qi et al. 2016), the Army Research Lab (Korolov et al. 2016), IARPA (Doyle et al. 2014; Hua, Lu, et al. 2013; Korkmaz et al. 2016) or the Department of Homeland Security (Korolov et al. 2016). This current surge in national security related civil unrest prediction research can be arguably understood as a continuation of longstanding Open Source Intelligence (OSINT) efforts (Schaurer and Störger 2013). They gained momentum in the US after the "Japanese attack on Pearl Harbor" and resulted in intelligence agencies collecting and analyzing foreign media. After the attacks of 9/11, these practices broadened to an "ever expanding universe of open sources" which ensued cooperation with companies and universities. The latter has been deemed preferential in many cases for state actors because it avoids conflicts with "profit-oriented players" and "a fertile ground for capturing expertise." In tandem, the trend towards social media surveillance in computer science arguably also picked up "in 2001 after the terrorist attacks of 9/11" (Reuter and Kaufhold 2018, p. 1), as "in the following years [..] sometimes summarized under the term crisis informatics, a variety of studies focusing on the use of ICT and social media before, during or after nearly every crisis and emergency has arisen." (p. 1). In turn, civil unrest prediction as a subject of research has also emerged in response to these heightened anxieties and techno-security culture (Weber and Kämpf 2020). The recent COVID crisis and its ensuing heavy use of prediction technologies to control the pandemic (Heimstädt et al. 2020) could further reinforce these trends towards evermore prediction technologies for security purposes.

I also encountered industry, academia, military, and law enforcement cooperation in this space. For instance, the Data to Decisions CRC, an Australian government research program, funded research into civil unrest prediction. The project was concerned with "solving big data challenges in the national security community" (Data to Decisions CRC n.d.). It led to the creation of spin-off companies such as Fivecast, an "anti-terrorism data startup" (Powell 2019) aiming to expand their operations across the globe. Fivecast uses "publicly available data from social media platforms to provide insights for workers in law enforcement, defense and national intelligence" (Powell 2019). Their system "enables automated monitoring of large sets of data to identify a wide range of threats, such as group violence, protest activity or lone actor activity" (Fivecast n.d.). This company thereby frames unrest, in this case, "protest activity," as a threat or security risk which should be anticipated. Its self-description as an anti-terrorism startup raises ethical questions around its involvement in protest surveillance as these two activities are not necessarily related. Another example is the IARPA-funded research project EMBERS, one of the most extensive civil unrest prediction research efforts I encountered. It was also connected to commercial actors since it was also based on an "industry-university partnership"(Ramakrishnan et al. 2014, p. 1). These cases highlight a network made of industry, academia, law enforcement, intelligence, and the military brought together by technology for predicting unrest activity to preempt (national) security threats. It shows further that research into civil unrest prediction is not simply neutral but implicated in a broader security and surveillance apparatus that sees unrest and protest as risky.

4.2. Civil unrest as economic risk

Civil unrest was also often framed as an economic risk, although most research projects were seemingly not industry-funded. In particular, the prediction of (labor) strikes was a commonly mentioned example for economic risk (Chen and Neill 2014; Hossny and Mitchell 2018; Hua, Lu, et al. 2013; Korkmaz et al. 2016; Muthiah et al. 2015; Qi et al. 2016; Xu et al. 2014; Zhao et al. 2017). This focus on labor organizing practices also relates civil unrest prediction to the labor union avoidance industry, which seeks to preempt unionization and organization to avoid ensuing costs, e.g., through worker benefits and wage rises. Researchers part of the EMBERS project also

discussed "labor and student unions" (Muthiah, Vullikanti, et al. 2016, p. 212-213) as prediction targets and stated that due to their size, they often also circulate announcements, which improves the quality of unrest predictions. In many papers, the prediction of labor organizing was just one in a list of possible applications, including the previously mentioned national security risks. This ascribed flexibility of the technology to different domains is concerning, as it potentially associates through risk calculations labor organizing activity and other forms of democratic protest possibly with the national security domain as they could also be or become risky to states.

Unrest was framed as a potentially risky disruption to flows of capital by researchers, e.g. because it can "cause disruptions to supply chain logistics, travel, and other sectors, and anticipating disruptions is key to ensuring safety as well as reliability"(Muthiah, Vullikanti, et al. 2016). The authors further strengthen this characterization of unrest as costly and risky by arguing that: "given the vulnerability of large gatherings to provocation by handfuls of violence-oriented protestors (e.g., Black Box anarchists in Brazil) the economic, social and political costs of large-scale public demonstrations are also potentially significant to marchers, bystanders, property owners and the government – democratically elected or not" (p. 213). Also, peaceful protests were framed as risky as they also supposedly could disrupt flows of capital: "There are economic costs to even peaceful disruptions embodied in civil unrest due to lost work hours and the deployment of police to manage traffic and the interactions between protestors and bystanders" (p. 213).

The presented 'fix' to such costly risks was the preemptive potential of civil unrest prediction as, for instance, "companies with personnel and supply chain operations can ask their employees to stay at home and to remain apolitical and can attempt to safeguard their facilities in advance" (Kallus 2014, p. 626). Researchers pointed to prediction as a remedy for unrest risk to tourism. It could "inform tourists of protest-prone zones" (Ansah et al. 2018a) or serve as a "guidance for travel planning" (Kang et al. 2017, p. 865) so that "users can gain insight into safety condition of different places by observing the distribution of reported and predicted civil unrest events" (p. 865). The risk of traffic disruptions was promised to be preempted by "help[ing] traffic regulators divert traffic effectively" (Ansah et al. 2018a). All these cases highlight the promise of civil unrest prediction to enable anticipatory governance of risky disruptions to flows of capital across various industries.

This interest of researches into civil unrest prediction as an economic risk assessment technology for companies can also be ascribed in part to the neoliberalization of universities (Canaan and Shumar 2008; Lave et al. 2010), which has entailed a "narrowing of research agendas to focus on the needs of commercial actors." I also encountered university spin-offs that illustrate a research to product pipeline. One example is the Austrian company Prewave which offers "realtime and predictive risk alerts" (Prewave n.d.) of supply chain disruptions such as labor strikes through a dashboard. It was based on a dissertation. One publication (Purwarianti et al. 2016), which one of the founders of Prewave has co-authored, introduced an information extraction system, which likely was foundational to the company. It retrieves event information from tweets on upcoming labor strikes in Indonesia to aid in their anticipation and detection. Similarly, Fivecast, which I mentioned in the previous section, did undergo this transformation from research project to civil unrest risk assessment technology company. I also encountered a consultancy that recommended the neoliberal university as a desirable partner for industries (John 2015). It argued that companies should actively reach out to universities about building supply chain analytics systems (IBM 2019), which often include unrest detection/prediction (NC4 n.d.), since they are often willing to cut the costs of resulting beta software.

The examples of civil unrest prediction research in industry also highlight neoliberal logics within contemporary academia. They show scholars engaged in industry partnerships and spin-off companies, which may explain in part the academic interest in unrest as an economic risk to be mitigated. This focus on corporate interests points to a need within academia to reflect and discuss whose risks are mitigated by current research efforts and what kinds of power relations are made durable (D'Ignazio and Klein 2020). Universities, especially publicly funded institutions, should seek to recenter worker needs and rights in their research practice. This point is made strongly clear by how civil unrest prediction systems are being developed under the guise of seemingly neutral research while their intended uses are also for the surveillance and prediction of labor strikes and unionization efforts. They are further described as tools to undermine such efforts. One Australian researcher has argued, for instance in an interview (Gibson 2018) on protest and strike prediction based on social media data, that the police can use this technology to "plan for disruptive events and hopefully divert them" and that the public can be warned about them ahead of time. Similarly, the above mentioned company Prewave has advertised their technology to companies as a "shitstorm insurance" (Grill 2020), framing possibly important grievances of activists and workers as mere public relations risks. In the introduction, I have also listed a few companies like Walmart and Wholefoods reportedly using risk assessment technologies based on various data to anticipate labor organizing. In particular, the parent company of Wholefoods, Amazon, has also been recently heavily critiqued for its anti-unionization tactics (McNicholas 2021). These examples illustrate how civil unrest prediction promises also to aid in averting various kinds of economic disruption and thereby potentially weakening possibilities for workers to make their grievances heard, felt, and resolved. It is thereby a technology with worrying implications for worker rights.

This section on civil unrest prediction as risk assessment has highlighted how civil unrest is framed in the analyzed publications as a (national) security risk and a risk to flows of capital. These imaginaries of unrest as a source of risk in need of surveillance and 'fixing' to ensure stability and security are co-constructed in tandem with prediction technologies by various actors such as governments, researchers, and companies. Ultimately, civil unrest prediction technology is often framed as a (national) security and economic risk assessment tool, which marks future civil unrest as risky and knowable. Various actors build a network around this technology stabilized through different interests such as funding incentives for researchers, profit for companies, and need for expertise by government agencies. Concerningly, the perspective that is mostly absent in this network belongs to the targeted protestors. In the next section, I unpack justifications of researchers for their work and discuss its political implications.

5. Justifications for civil unrest prediction research

Most of the civil unrest prediction research I analyzed targets all kinds of protests and thereby also marks them as risky, which could result in various detrimental treatments such as preemptive interventions, policing, increased surveilling, and targeting of individuals and groups. The described users of civil unrest prediction are usually in positions of power, such as law enforcement officials or big companies, and seemingly aim to employ the technology to mitigate and preempt protest before it becomes risky or disruptive. The thresholds for when protests become too risky are unclear. Knowledge is always "produced with a particular interest" (Baaz et al. 2017, p. 139). In turn, unrest prediction should not be understood as somehow neutral. Unrest prediction is developed in a moment of anxieties over fears of the disruption of the social order and challenges to the legitimacy of powerful actors (Solnit 2010). It embodies certain politics and interests, which I have tried to highlight in particular in section four. There is a need for debate around power, justice, and how they are inscribed (Akrich 1992) into research and uses of civil unrest prediction. In this section, I unpack

how scholars in this field have justified their work in various ways and highlight how these elaborations miss essential perspectives.

The researchers have motivated civil unrest prediction and detection by highlighting how it could mitigate various risks. Some researchers also argued for the potential of the technology to facilitate communication and understanding between protestors and decision makers. One team of authors highlighted how civil unrest prediction has the potential to help "service providers prioritize on the concerns of citizens" (Ansah et al. 2018a) as it informs them before unrest follows. The analysis of identified and predicted unrest events was also argued to "provide new insights to authorities and policymakers [on how] to understand issues of public unrest, and to identify opinions and expressions on a sensitive topic like race, at a scale and scope not possible through conventional means such as surveys"(De Choudhury et al. 2016, p. 100). Furthermore, an "effective protest forecasting system" (Ramakrishnan et al. 2014, p. 1800) was presented as able to "contribute to making the transmission of citizen preferences to government less costly to the economy and society, by enabling governments to respond to high priority grievances in advance of anticipated protests. If the response by the government causes a cancellation or lower turnout for the event, this decreases the costs incurred by even peaceful disruptions" (p. 1800). Put differently, civil unrest prediction also promises to facilitate a form of early communication that could lead to earlier diversion of unrest as priorities of protestors are met.

This supposed benefit of early communication of grievances mischaracterizes unrest as emerging out of a communication problem in contrast to the building of pressure and counter-power to elicit change. It is not unlikely that authorities do not agree with protestors' motives, as the voicing of grievances usually comes before unrest when demands are not met. In turn, framing civil unrest prediction simply as a technique for early communication misses its potential to divert or preempt protest through various means other than meeting priorities and wishes. The technology could, in turn, also become a potential danger for protestors as it's used to track and target them. Furthermore, if the goal indeed was early communication, then such a surveillance-based approach invisible to protestors is misguided. First, it presumes that only at the point when masses mobilize should grievances be met. Second, activists are often willing to communicate their grievances and requests to different actors. Statistical social media analysis from which supposed preferences are inferred or guessed is not superior to direct communication. Ultimately, instead of error-prone prediction systems, stable new communication structures with direct lines to governments and companies to voice grievances could be built up. Especially in global supply chains, infrastructure to control suppliers and ensure global fair working standards are needed to improve the current situation of workers. This requires a reimagining of existing social and legal structures and cannot be 'fixed' with simple social media surveillance tools. In conclusion, this early communication frame is a problematic justification for this kind of work. Furthermore, the framing of peaceful protest as costly and risky by government-funded researchers, like the EMBERS team, is deeply problematic in itself. The protestors are exercising a democratic right, and the mentioned economic costs certainly do not outweigh it. The quote points to a foundational tension at the heart of civil unrest prediction, which is about what unrest activity is an acceptable target for this kind of surveillance and the issues of misuse related to, e.g., averting democratic and emancipatory protests.

Another justification of civil unrest prediction was the supposed benefit to science for better understanding the social world. One researcher argued, "socio-technical advances have created favorable conditions for forecasting certain types of mobilizations or protests while simultaneously generating large reservoirs of online data" (Qi et al. 2016, p. 2). They further offer "a new avenue of research where information flows may provide fresh insight into human behavior" (p. 2). These

statements frame the current internet landscape as a ripe, unexplored land with opportunity which should be extracted in the name of science to improve scholarly knowledge. This points to an underlying "colonial impulse" (Dourish and Mainwaring 2012) of civil unrest prediction research as it aims to colonize the seemingly unknown online world to harvest new knowledge. However, in this quest, activists' and protestors' perspectives and possible concerns are sidelined in favor of supposed discovery. It highlights a positivist logic that frames science and knowledge as neutral ends that should simply be pursued and increased. A statement of the EMBERS researchers also illustrates this. They argued that "the appropriate safeguards require developing transparent and accountable democratic systems, not outlawing science" (Muthiah, Vullikanti, et al. 2016) and that the "potential power of civil unrest forecasting systems, like those of most scientific advances, is susceptible to abuse by both democratic and non-democratic governments" (p. 213). The authors thereby act as if science is and should be somehow detached from society as only uses need to be regulated, but decades of research in STS have shown this isn't the case since science and technology in themselves also embody politics (Parthasarathy 2015; Sarewitz 2011; Winner 1980). The logic of the authors justifies almost any research, no matter its potential harms. It highlights a need for reflection and engagement of big data researchers with activists and fields such as STS, specifically for those in civil unrest prediction. The statements further frame civil unrest prediction as scientific advancement, but this understanding is likely not shared by those surveilled and requires democratic deliberation. Even if well-intended, research can be harmful to certain people and contribute to their marginalization.

One of the most central justifications of civil unrest prediction was the public availability of online data, which was described as having received "unprecedented attention over the past decade" (Kang et al. 2017, p. 1) due to its "ubiquity" (p. 1). One author, for instance, argued that content posted to social media is "public and accessible as open source data" (Agarwal 2017, p. 2). Nevertheless, just because data is publicly collectible does not mean it should be mined, or even that social media users want or would be okay with being analyzed for purposes they may disagree with. One study (Fiesler and Proferes 2018) found, for instance, that a majority of social media users are not aware of researchers using tweets in their work and think they should not do so without permission. Ultimately, this simple dichotomy between private and public is problematic. Just because certain activities are done in public, such as posting on social media, does not mean they should not be protected (Chun 2016). This justification centered simply on the ascribed publicness of data frames research built on top of it as an unpolitical endeavor. The ascribed publicness of data ultimately depoliticizes it. The justification indicates a stance that holds anything goes when data is somehow accessible. It is important to note that even if people on social media give their consent, research should also aim to be beneficial, especially to marginalized communities, which this technology, in many cases, may not be. Most publications concerningly did not discuss negative consequences to protestors through this technology and its varied uses. My analysis thereby highlights a critical need for discussion and reflection in this area of work.

6. Conclusion

In this work, I have first attempted to conceptualize civil unrest prediction. I showed its temporal entwinement with detection and argued that it adheres to the big data paradigm by highlighting data sources, features, and methods it is based upon. Then I showed how researchers frame unrest as both national security and economic risk in need of 'fixing' or intervention. In turn, civil unrest prediction technologies are also presented as risk assessment tools that mark future unrest as risky and calculable. Finally, I unpacked justifications provided by researchers for conducting this kind of work.

My analysis shows a critical need for reflection and critical research on academic practices at the intersection of data science, risk assessment, and unrest prediction, as well as on the ethics and politics of protest research and ensuing technological applications. This need is further strengthened by the COVID-19 pandemic, increasing economic inequality, and the climate change crisis, which may lead to increases in unrest activity over the next years. These trends could likely result in an expansion of risk assessment and surveillance capabilities of governments and companies. Academia, especially computer science, will have a role in this development and needs to take a stance as the technologies it is responsible for are potentially harmful to marginalized communities. Researchers should engage more with underlying political questions of civil unrest prediction, such as if this work should be done at all. Moreover, if the technology seems appropriate for certain use cases, how can it be ensured that these uses are legitimate and conducted with transparency, justice, and accountability in mind. There are potential legitimate uses of civil unrest prediction, e.g., for uncovering potential human rights abuses or shedding light on unrest across the globe that seeks more visibility as it faces suppression. Still, the question if these supposed benefits are enough to motivate further research and in what forms remains. There are risks to various people, democracy, and activists that may arise through the use of such mass surveillance technology, such as preemptive interventions against protest activity like potential misinformation campaigns targeted at just emerging protest activity. These dangers also need to be curbed. It is important to note that as long as social media remains available as a data source for such tools, certain actors may seek to exploit it for their ends. In turn, various forms of social media analysis by researchers could also be seen as a potential mechanism to add some accountability to platforms otherwise only accessible to those with the appropriate knowledge and resources. However, such work should be grounded in actual needs of marginalized communities (Costanza-Chock 2020).

The structural and historical inequities prevalent in contemporary societies, such as racism, over-policing, and rampant income inequality, show how civil unrest prediction should also be understood as part of the "New Jim Code" (Benjamin 2019). This concept describes and problematizes technology presented as objective and neutral while reinforcing racist politics. Civil unrest prediction is a technology that further enables the surveillance of resistance movements, like the Black Lives Matter protests, and in turn, potentially undermines efforts to challenge structural racism. Consequently, I argue for a stop of this kind of research and its use by governmental agencies. Furthermore, an evaluation of its impacts and aims is needed. This call also extends to companies that employ civil unrest prediction in any capacity. My analysis has highlighted how the technology is also understood as a tool for the preemption of labor organizing and, in turn, potentially reinforcing inequalities and unfair labor practices. Thus, there is an urgent need for new regulation to strengthen human and worker rights across global supply chains and collective global organizing to implement accountability mechanisms and remedy exploitative practices. A leaked draft on new EU regulation for AI promisingly mentioned the prediction of "events of social unrest" (Wagner 2021) as a high-risk application. This classification would have required deployed systems to adhere to strong transparency and accountability rules. However, this mention has seemingly been removed later in the official proposal (European Commission 2021), which raises questions in what capacity in the EU, and other jurisdictions possibly inspired by its regulations, such systems will be scrutinized considering all the risks they pose to various people. In general, this space of unrest prediction based on publicly accessible data is currently not strongly regulated, and pioneering laws centering the perspectives of marginalized communities are needed. It is also important to note that the analysis of characteristics of past protests is a separate matter. Still, it also comes with a set of related ethical questions, such as should social media data be used just because it can be somehow accessed.

I have three recommendations for future data science studies concerned with civil unrest. First, I ask researchers working in the intersection of social movement research and data science to provide

statements on ethical and political considerations of their work, and especially also to ask questions around for whom and for what purpose they are producing knowledge. Second, research should be more concerned with the politics and technical limitations of this project. Similar to how currently researchers almost only cite seemingly successful related research projects, they also should engage with limitations, social implications, and potential harms to various groups of people. There is a great amount of literature on the politics and limitations of big data, seemingly ignored. Third, a deeper engagement and cooperation with the social sciences and marginalized communities are essential to ensure that research is beneficial and not harmful. I aim to deconstruct promises associated with the technology and further illustrate its technical affordances and embedded politics in future work.

References

Agarwal, Swati (2017). Open source social media intelligence for enabling government applications: extended abstract. *ACM SIGWEB Newsletter*, no. Summer, July 2017, pp. 1–19.

Agarwal, Swati; and Ashish Sureka (2015). Applying social media intelligence for predicting and identifying on-line radicalization and civil unrest oriented threats. *arXiv preprint arXiv:1511.06858*, 2015.

Agarwal, Swati; and Ashish Sureka (2016). Investigating the Potential of Aggregated Tweets as Surrogate Data for Forecasting Civil Protests. In *CODS '16. Proceedings of the 3rd IKDD Conference on Data Science, 2016*. Pune, India: ACM Press, pp. 1–6.

Agarwal, Swati; Ashish Sureka; and Vikram Goyal (2015). Open source social media analytics for intelligence and security informatics applications. In *International Conference on Big Data Analytics*. Springer, pp. 21–37.

Ahmed, Nafeez (2018). Pentagon Wants to Predict Anti-Trump Protests Using Social Media Surveillance. *Vice*. https://www.vice.com/en_us/article/7x3g4x/pentagon-wants-to-predict-anti-trump-protests-using-social-media-surveillance. Accessed 15 May 2020

Akrich, Madeleine (1992). The de-scription of technical objects, 1992.

Alsaedi, Nasser; Pete Burnap; and Omer Rana (2017). Can We Predict a Riot? Disruptive Event Detection Using Twitter. *ACM Transactions on Internet Technology*, vol. 17, no. 2, March 2017, pp. 1–26.

Amoore, Louise (2020). *Cloud ethics: algorithms and the attributes of ourselves and others*. Durham: Duke University Press.

Ansah, Jeffery; Wei Kang; Lin Liu; Jixue Liu; and Jiuyong Li (2018a). Information propagation trees for protest event prediction. In *PAKDD 2018. Pacific-Asia Conference on Knowledge Discovery and Data Mining*. Springer, pp. 777–789.

Ansah, Jeffery; Wei Kang; Lin Liu; Jixue Liu; and Jiuyong Li (2018b). SensorTree: Bursty Propagation Trees as Sensors for Protest Event Detection. In *International Conference on Web Information Systems Engineering*. Springer, pp. 281–296.

Avle, Seyram; Cindy Lin; Jean Hardy; and Silvia Lindtner (2020). Scaling Techno-Optimistic Visions. *Engaging Science, Technology, and Society*, vol. 6, pp. 237–254.

Baaz, Mikael; Mona Lilja; and Stellan Vinthagen (2017). *Researching Resistance and Social Change: A Critical Approach to Theory and Practice*. London ; New York: Rowman & Littlefield Publishers.

Bahrami, Mohsen; Yasin Findik; Burcin Bozkaya; and Selim Balcisoy (2018). Twitter Reveals: Using Twitter Analytics to Predict Public Protests. *arXiv:1805.00358 [cs]*, May 2018. http://arxiv.org/abs/1805.00358. Accessed 15 March 2019

Beck, Ulrich (1992). *Risk society: Towards a new modernity*. Sage.

Benjamin, Ruha (2019). *Race after technology: Abolitionist tools for the new jim code*. John Wiley & Sons.

Benkhelifa, Elhadj; Elliott Rowe; Robert Kinmond; Oluwasegun A Adedugbe; and Thomas Welsh (2014). Exploiting social networks for the prediction of social and civil unrest: A cloud based framework. In *2014 international conference on future internet of things and cloud*. IEEE, pp. 565–572.

Borradaile, Glencora; Brett Burkhardt; and Alexandria LeClerc (2020). Whose tweets are surveilled for the police: an audit of a social-media monitoring tool via log files. In *FAT* '20. Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. Barcelona, Spain: ACM, pp. 570–580.

Bowen, Glenn A (2009). Document Analysis as a Qualitative Research Method. *Qualitative Research Journal*, vol. 9, no. 2, August 2009, pp. 27–40.

boyd, danah (2010). Social network sites as networked publics: Affordances, dynamics, and implications. In *A networked self*. Routledge.

boyd, danah; and Kate Crawford (2012). Critical questions for big data: Provocations for a cultural, technological, and scholarly phenomenon. *Information, communication & society*, vol. 15, no. 5, 2012, pp. 662–679.

Cadena, Jose; Gizem Korkmaz; Chris J. Kuhlman; Achla Marathe; Naren Ramakrishnan; and Anil Vullikanti (2015). Forecasting Social Unrest Using Activity Cascades. *PLOS ONE*, vol. 10, no. 6, June 2015, pp. e0128879.

Canaan, Joyce E.; and Wesley Shumar (2008). *Structure and Agency in the Neoliberal University*. Routledge.

Cederman, Lars-Erik; and Nils B. Weidmann (2017). Predicting armed conflict: Time to adjust our expectations? *Science*, vol. 355, no. 6324, February 2017, pp. 474–476.

Chen, Feng; and Daniel B. Neill (2014). Non-parametric scan statistics for event detection and forecasting in heterogeneous social media graphs. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, pp. 1166–1175.

Chen, Feng; and Daniel B Neill (2015). Human Rights Event Detection from Heterogeneous Social Media Graphs. *Big Data*, vol. 3, no. 1, March 2015, pp. 34–40.

Chun, Wendy Hui Kyong (2016). *Updating to Remain the Same: Habitual New Media*. Cambridge, MA: The MIT Press.

Clarke, Adele E.; Carrie Friese; and Rachel S. Washburn (2017). *Situational analysis: grounded theory after the interpretive turn*. Sage Publications.

Colbaugh, R.; and K. Glass (2010). Early warning analysis for social diffusion events. In *2010 IEEE International Conference on Intelligence and Security Informatics*. IEEE, pp. 37–42.

Colbran, Marianne P (2018). Policing, social media and the new media landscape: can the police and the traditional media ever successfully bypass each other? *Policing and Society*. pp. 1–15.

Compton, Ryan; Craig Lee; Tsai-Ching Lu; Lalindra De Silva; and Michael Macy (2013). Detecting future social unrest in unprocessed Twitter data: "Emerging phenomena and big data." In *2013 IEEE International Conference on Intelligence and Security Informatics*. Seattle, WA, USA: IEEE, pp. 56–60.

Compton, Ryan; Craig Lee; Jiejun Xu; Luis Artieda-Moncada; Tsai-Ching Lu; Lalindra De Silva; and Michael Macy (2014). Using publicly visible social media to build detailed forecasts of civil unrest. *Security Informatics*, vol. 3, no. 1, September 2014, pp. 4.

Costanza-Chock, Sasha (2020). *Design Justice: Community-Led Practices to Build the Worlds We Need*. MIT Press.

Data to Decisions CRC (n.d.). D2D Legacy Website. https://www.d2dcrc.com.au/article-content/d2d-crc-spinout-company-nqry. Accessed 14 June 2020

De Choudhury, Munmun; Shagun Jhaver; Benjamin Sugar; and Ingmar Weber (2016). Social media participation in an activist movement for racial equality. In *Tenth International AAAI Conference on Web and Social Media*.

Dencik, Lina; Arne Hintz; and Zoe Carey (2018). Prediction, pre-emption and limits to dissent: Social media and big data uses for policing protests in the United Kingdom. *New Media & Society*, vol. 20, no. 4, April 2018, pp. 1433–1450.

D'Ignazio, Catherine; and Lauren F. Klein (2020). *Data Feminism*. Cambridge, Massachusetts: The MIT Press.

Dilanian, Ken (2021). DHS is gathering intelligence on security threats from social media. *NBC News*. https://www.nbcnews.com/politics/national-security/dhs-launches-warning-system-find-domestic-terrorism-threats-public-social-n1266707. Accessed 13 May 2021

Dourish, Paul; and Scott D. Mainwaring (2012). Ubicomp's Colonial Impulse. In *UbiComp '12. Proceedings of the 2012 ACM Conference on Ubiquitous Computing*. New York, NY, USA: ACM, pp. 133–142.

Doyle, Andy; Graham Katz; Kristen Summers; Chris Ackermann; Ilya Zavorin; Zunsik Lim; et al (2014). The EMBERS architecture for streaming predictive analytics. In *2014 IEEE international conference on big data (big data)*. IEEE, pp. 11–13.

Egawhary, Elena M (2019). The Surveillance Dimensions of the Use of Social Media by UK Police Forces. *Surveillance & Society*, vol. 17, no. 1/2, March 2019, pp. 89–104.

Ensign, Danielle; Sorelle A. Friedler; Scott Neville; Carlos Scheidegger; and Suresh Venkatasubramanian (2017). Runaway feedback loops in predictive policing. *arXiv preprint arXiv:1706.09847*, 2017.

Ertugrul, Ali Mert; Yu-Ru Lin; Wen-Ting Chung; Muheng Yan; and Ang Li (2019). Activism via attention: interpretable spatiotemporal learning to forecast protest activities. *EPJ Data Science*, vol. 8, no. 1, December 2019, pp. 5.

European Commission (2021). Proposal for a Regulation laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). https://digital-strategy.ec.europa.eu/en/library/proposal-regulation-laying-down-harmonised-rules-artificial-intelligence. Accessed 31 May 2021

Ferguson, Andrew Guthrie (2017). *The Rise of Big Data Policing: Surveillance, Race, and the Future of Law Enforcement*. NYU Press.

Fiesler, Casey; and Nicholas Proferes (2018). "Participant" Perceptions of Twitter Research Ethics. *Social Media+ Society*, vol. 4, no. 1, 2018, pp. 2056305118763366.

Filchenkov, Andrey A.; Artur A. Azarov; and Maxim V. Abramov (2014). What is more predictable in social media: Election outcome or protest action? In *EGOSE '14. Proceedings of the 2014 Conference on Electronic Governance and Open Society: Challenges in Eurasia*. St. Petersburg, Russian Federation: ACM Press, pp. 157–161.

Fivecast (n.d.). Solutions. *Fivecast*. https://www.fivecast.com/solutions/. Accessed 14 June 2020

Ganar, Ruchika Parmarth; and Shrikant Ardhapurkar (2016). Prediction of civil unrest by analysing social network using keyword filtering: A survey. In *2016 Online International Conference on Green Engineering and Technologies (IC-GET)*. IEEE, pp. 1–4.

Gandy, Oscar H (2016). *Coming to terms with chance: Engaging rational discrimination and cumulative disadvantage*. Routledge.

Gee, James Paul (2014). *An introduction to discourse analysis: Theory and method*. Routledge.

Gerber, Matthew S (2014). Predicting crime using Twitter and kernel density estimation. *Decision Support Systems*, vol. 61, 2014, pp. 115–125.

Gibson, Candy (2018). How social media can predict protests and strikes. https://web.archive.org/web/20200415234752/http://w3.unisa.edu.au/unisanews/2018/March/feature.asp. Accessed 9 April 2019

Goode, Brian J.; Siddharth Krishnan; Michael Roan; and Naren Ramakrishnan (2015). Pricing a Protest: Forecasting the Dynamics of Civil Unrest Activity in Social Media. *PLOS ONE*, vol. 10, no. 10, October 2015, pp. e0139911.

Grill, Gabriel (2020). RESEARCHING ONLINE LABOR STRIKE AND PROTEST PREDICTION TECHNOLOGIES. *AoIR Selected Papers of Internet Research*, 2020.

Gupta, Dipak K; Sathappan Muthiah; David Mares; and Naren Ramakrishnan (2017). Forecasting Civil Strife: An Emerging Methodology. *The Third International Conference on Human and Social Analytics*, 2017, pp. 6.

Hälterlein, Jens; and Lars Ostermeier (2018). Special Issue: Predictive Security Technologies. *European Journal for Security Research*, vol. 3, no. 2, October 2018, pp. 91–94.

Hargittai, Eszter (2015). Is bigger always better? Potential biases of big data derived from social network sites. *The ANNALS of the American Academy of Political and Social Science*, vol. 659, no. 1, 2015, pp. 63–76.

Heimstädt, Maximilian; and Leonhard Dobusch (2021). Riskante Retweets: „Predictive Risk Intelligence" und Interessenvertretung in globalen Wertschöpfungsnetzwerken. I*ndustrielle Beziehungen,* vol. 28, no. 2.

Heimstädt, Maximilian; Simon Egbert; and Elena Esposito (2020). A Pandemic of Prediction: On the Circulation of Contagion Models between Public Health and Public Safety. *Sociologica*, vol. 14, no. 3, 2020, pp. 1–24.

Hermida, Alberto; and Víctor Hernández-Santaolalla (2018). Twitter and video activism as tools for counter-surveillance: the case of social protests in Spain. *Information, Communication & Society*, vol. 21, no. 3, March 2018, pp. 416–433.

Hoegh, Andrew; Scotland Leman; Parang Saraf; and Naren Ramakrishnan (2015). Bayesian Model Fusion for Forecasting Civil Unrest. *Technometrics*, vol. 57, no. 3, July 2015, pp. 332–340.

Hoffmann, Anna Lauren; Nicholas Proferes; and Michael Zimmer (2018). "Making the world more open and connected": Mark Zuckerberg and the discursive construction of Facebook and its users. *New Media & Society*, vol. 20, no. 1, January 2018, pp. 199–218.

Hossny, Ahmad H.; and Lewis Mitchell (2018). Event Detection in Twitter: A Keyword Volume Approach. In *2018 IEEE International Conference on Data Mining Workshops (ICDMW)*. IEEE, pp. 1200–1208.

Hsieh, Hsiu-Fang; and Sarah E. Shannon (2005). Three approaches to qualitative content analysis. *Qualitative health research*, vol. 15, no. 9, 2005, pp. 1277–1288.

Hua, Ting; Feng Chen; Liang Zhao; Chang-Tien Lu; and Naren Ramakrishnan (2013). STED: semi-supervised targeted-interest event detectionin in twitter. In *KDD '13. Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*. Chicago, Illinois, USA: ACM Press, pp. 1466.

Hua, Ting; Chang-Tien Lu; Naren Ramakrishnan; Feng Chen; Jaime Arredondo; David Mares; and Kristen Summers (2013). Analyzing civil unrest through social media. *Computer*, vol. 46, no. 12, 2013, pp. 80–84.

IARPA (2011). OSI. *Open Source Indicators (OSI)*. https://www.iarpa.gov/index.php/research-programs/osi?highlight=WyJvc2kiXQ==. Accessed 17 June 2020

IBM (2019). What is supply chain analytics? https://www.ibm.com/uk-en/supply-chain/supply-chain-analytics. Accessed 29 June 2020

Jasanoff, Sheila (2004). The idiom of co-production. *States of knowledge: The co-production of science and social order*, 2004, pp. 1–12.

John, Geraint (2015). The Changing Face of Supply Chain Risk Management. SCM World. https://web.archive.org/web/20200904024412/scmworld.com/wp-content/uploads/2017/01/The-Changing-Face-of-Supply-Chain-Risk-Management.pdf

Kallus, Nathan (2014). Predicting Crowd Behavior with Big Public Data. In *WWW '14 Companion. Proceedings of the 23rd International Conference on World Wide Web.* Seoul, Korea: ACM, pp. 625–630.

Kang, Wei; Jie Chen; Jiuyong Li; Jixue Liu; Lin Liu; Grant Osborne; et al (2017). Carbon: forecasting civil unrest events by monitoring news and social media. In *International Conference on Advanced Data Mining and Applications*. Springer, pp. 859–865.

Keyes, Os (2018). The Misgendering Machines: Trans/HCI Implications of Automatic Gender Recognition. In *Proceedings ACM Human-Computer Interaction*, vol. 2, no. CSCW, November 2018, pp. 88:1-88:22.

Kitchin, Rob (2014). The real-time city? Big data and smart urbanism. *GeoJournal*, vol. 79, no. 1, February 2014, pp. 1–14.

Kline, Ronald; and Trevor Pinch (1996). Users as agents of technological change: The social construction of the automobile in the rural United States. *Technology and culture*, vol. 37, no. 4, 1996, pp. 763–795.

Korkmaz, Gizem; Jose Cadena; Chris J. Kuhlman; Achla Marathe; Anil Vullikanti; and Naren Ramakrishnan (2015). Combining Heterogeneous Data Sources for Civil Unrest Forecasting. In *ASONAM '15. Proceedings of the 2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2015*. Paris, France: ACM Press, pp. 258–265.

Korkmaz, Gizem; Jose Cadena; Chris J. Kuhlman; Achla Marathe; Anil Vullikanti; and Naren Ramakrishnan (2016). Multi-source models for civil unrest forecasting. *Social Network Analysis and Mining*, vol. 6, no. 1, December 2016.

Korolov, Rostyslav; Di Lu; Jingjing Wang; Guangyu Zhou; Claire Bonial; Clare Voss; et al (2016). On predicting social unrest using social media. In *2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*. San Francisco, CA, USA: IEEE, pp. 89–95.

Lave, Rebecca; Philip Mirowski; and Samuel Randalls (2010). Introduction: STS and Neoliberal Science: *Social Studies of Science*, September 2010.

Lazer, David; Ryan Kennedy; Gary King; and Alessandro Vespignani (2014). The parable of Google Flu: traps in big data analysis. *Science*, vol. 343, no. 6176, 2014, pp. 1203–1205.

Liu, Yu; Baojian Zhou; Feng Chen; and David W. Cheung (2016). Graph Topic Scan Statistic for Spatial Event Detection. In *CIKM '16. Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*. Indianapolis, Indiana, USA: ACM Press, pp. 489–498.

Luhmann, Niklas (2005). *Risk*. New Brunswick, N.J: Routledge.

Manrique, Pedro; Hong Qi; Ana Morgenstern; Nicolas Velásquez; Tsai-Ching Lu; and Neil Johnson (2013). Context matters: Improving the uses of big data for forecasting civil unrest: Emerging phenomena and big data. In *2013 IEEE International Conference on Intelligence and Security Informatics*. IEEE, pp. 169–172.

Marcus, George E.; and Michael MJ Fischer (2014). *Anthropology as cultural critique: An experimental moment in the human sciences*. University of Chicago Press.

Marwick, Alice E.; and Robyn Caplan (2018). Drinking male tears: language, the manosphere, and networked harassment. *Feminist Media Studies*, vol. 18, no. 4, July 2018, pp. 543–559.

McNicholas, Celine (2021). Amazon's anti-union campaign is part of a long history of employer opposition to organizing: Passing the PRO Act would be a critical first step. *Economic Policy Institute*. https://www.epi.org/blog/amazons-anti-union-campaign-is-part-of-a-long-history-of-employer-opposition-to-organizing-passing-the-pro-act-would-be-a-critical-first-step/. Accessed 31 May 2021

McQuillan, Dan (2018). Data Science as Machinic Neoplatonism. *Philosophy & Technology*, vol. 31, no. 2, June 2018, pp. 253–272.

Mishler, A.; K. Wonus; W. Chambers; and M. Bloodgood (2017). Filtering Tweets for Social Unrest. In *2017 IEEE 11th International Conference on Semantic Computing (ICSC)*. IEEE, pp. 17–23.

Mullet, Dianna R (2018). A General Critical Discourse Analysis Framework for Educational Research. *Journal of Advanced Academics*, vol. 29, no. 2, May 2018, pp. 116–142.

Muthiah, Sathappan; Bert Huang; Jaime Arredondo; David Mares; Lise Getoor; Graham Katz; and Naren Ramakrishnan (2015). Planned Protest Modeling in News and Social Media. In *AAAI*. pp. 3920–3927.

Muthiah, Sathappan; Bert Huang; Jaime Arredondo; David Mares; Lise Getoor; Graham Katz; and Naren Ramakrishnan (2016). Capturing Planned Protests from Open Source Indicators. *AI Magazine*, vol. 37, no. 2, July 2016, pp. 63–75.

Muthiah, Sathappan; Anil Vullikanti; Achla Marathe; Kristen Summers; Graham Katz; Andy Doyle; et al (2016). EMBERS at 4 years: Experiences operating an Open Source Indicators Forecasting System. In *KDD '16. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. San Francisco, California, USA: ACM Press, pp. 205–214.

NC4 (n.d.). Risk Management | NC4 | Business Continuity | Corporate Security | NC4. https://web.archive.org/web/20200624173922/https://www.nc4.com/risk. Accessed 24 June 2020

Osborne, Grant; Nick Lothian; Grant Neale; Terry Moscou; Andrew Nguyen; Jie Chen; et al (2017). The beat the news system: Forecasting social disruption via modelling of online behaviours. *Journal of the Australian Institute of Professional Intelligence Officers*, vol. 25, no. 2, 2017, pp. 35.

Parthasarathy, Shobita (2015). Governance lessons for CRISPR/Cas9 from the missed opportunities of Asilomar. *Ethics in Biology, Engineering and Medicine: An International Journal*, vol. 6, no. 3–4, 2015.

Pasquale, Frank (2015). *The black box society: The secret algorithms that control money and information*. Harvard University Press.

Pelzer, Robert (2018). Policing of Terrorism Using Data from Social Media. *European Journal for Security Research*, vol. 3, no. 2, October 2018, pp. 163–179.

Peterson, Hayley (2015). How Walmart Keeps an Eye on Its Massive Workforce. *Bloomberg.com*. http://www.bloomberg.com/features/2015-walmart-union-surveillance/. Accessed 13 May 2021

Peterson, Hayley (2020). Whole Foods tracks unionization risk with heat map. *Business Insider*. https://web.archive.org/web/20200422015312/https://www.businessinsider.com/whole-foods-tracks-unionization-risk-with-heat-map-2020-1?r=DE&IR=T. Accessed 26 April 2020

Pinch, Trevor; and Wiebe Bijker (1987). The Social Construction of Facts and Artifacts: Or How the Sociology of Science and the Sociology of Technology Might benefit Each Other. In *The Social Construction of Technological Systems*. Cambridge: MIT Press.

Powell, Dominic (2019, July 2). Anti-terrorism data startup Fivecast raises $4 million to crack US defence market. *SmartCompany*. https://www.smartcompany.com.au/startupsmart/news/anti-terrorism-defence-startup-fivecast/. Accessed 14 June 2020

Prewave (n.d.). Prewave | Artificial Intelligence Meets Risk Intelligence. *Prewave*. https://www.prewave.ai. Accessed 14 August 2019

Purwarianti, Ayu; Lisa Madlberger; and Mochammad Ibrahim (2016). Supervised Entity Tagger for Indonesian Labor Strike Tweets using Oversampling Technique and Low Resource Features. *TELKOMNIKA (Telecommunication Computing Electronics and Control)*, vol. 14, no. 4, 2016, pp. 1462–1471.

Qi, Hong; Pedro Manrique; Daniela Johnson; Elvira Restrepo; and Neil F. Johnson (2016). Open source data reveals connection between online and on-street protest activity. *EPJ Data Science*, vol. 5, no. 1, May 2016, pp. 18.

Ramakrishnan, Naren; Gizem Korkmaz; Chris Kuhlman; Achla Marathe; Liang Zhao; Ting Hua; et al (2014). "Beating the news" with EMBERS: forecasting civil unrest using open source indicators. In *KDD '14. Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. New York, New York, USA: ACM Press, pp. 1799–1808.

Ramakrishnan, Naren; Chang-Tien Lu; Madhav Marathe; Achla Marathe; Anil Vullikanti; Stephen Eubank; et al (2015). Model-Based Forecasting of Significant Societal Events. *IEEE Intelligent Systems*, vol. 30, no. 5, September 2015, pp. 86–90.

Renaud, Molly; Rostyslav Korolov; David Mendonça; and William Wallace (2019). Social Network Structure as a Predictor of Social Behavior: The Case of Protest in the 2016 US Presidential Election. In *XVIII International Conference on Data Science and Intelligent Analysis of Information*. Cham: Springer International Publishing, pp. 267–278.

Reuter, Christian; and Marc-André Kaufhold (2018). Fifteen years of social media in emergencies: A retrospective review and future directions for crisis Informatics. *Journal of Contingencies and Crisis Management*, vol. 26, no. 1, 2018, pp. 41–57.

Richardson, Rashida; Jason Schultz; and Kate Crawford (2019). Dirty Data, Bad Predictions: How Civil Rights Violations Impact Police Data, Predictive Policing Systems, and Justice. *NYUL Rev. Online*, vol. 94, 2019, pp. 15-55.

Rule, William; Wanyi Duan; Nitesh Prakash; Ni Zhuang; Rafael C Alvarado; and Donald E Brown (2018). Social pressure analysis of local events using social media data. In *2018 systems and information engineering design symposium (SIEDS)*. IEEE, pp. 277–281.

Sarewitz, Daniel (2011). *Institutional ecology and the social outcomes of scientific research*. Stanford University Press Stanford, CA.

Schaurer, Florian; and Jan Störger (2013). The evolution of open source intelligence (OSINT). *Comput Hum Behav*, vol. 19, 2013, pp. 53–56.

Singh, Sulbha; and Rajeev Pal (2018). Characterizing and Detecting Social Outrage on Twitter: Patel Reservation in Gujarat. In B. Panda, S. Sharma, and N. R. Roy (Eds.), *Communications in Computer and Information Science. Data Science and Analytics*. Singapore: Springer, pp. 498–515.

Solnit, Rebecca (2010). *A Paradise Built in Hell: The Extraordinary Communities That Arise in Disaster*. New York: Penguin.

Su, Chuan-Jun; and Yin-An Chen (2018). Risk assessment for global supplier selection using text mining. *Computers & Electrical Engineering*, vol. 68, May 2018, pp. 140–155.

Tan, Li; Suma Ponnam; Patrick Gillham; Bob Edwards; and Erik Johnson (2013). Analyzing the impact of social media on social movements: a computational study on Twitter and the Occupy Wall Street movement. In *Proceedings of the 2013 IEEE/ACM international conference on advances in social networks analysis and mining*. pp. 1259–1266.

Tufekci, Zeynep (2014). Big questions for social media big data: Representativeness, validity and other methodological pitfalls. *arXiv preprint arXiv:1403.7400*, 2014.

Tufekci, Zeynep (2017). *Twitter and tear gas: The power and fragility of networked protest*. Yale University Press.

Tufekci, Zeynep; and Christopher Wilson (2012). Social Media and the Decision to Participate in Political Protest: Observations From Tahrir Square. *Journal of Communication*, vol. 62, no. 2, April 2012, pp. 363–379.

Uldam, Julie (2018). Between visibility and surveillance: Challenges to anti-corporate activism in social media. In *Protests in the Information Age: Social Movements, Digital Practices and Surveillance*. London: Routledge.

Valenzuela, Sebastián; Arturo Arriagada; and Andrés Scherman (2012). The Social Media Basis of Youth Protest Behavior: The Case of Chile. *Journal of Communication*, vol. 62, no. 2, April 2012, pp. 299–314.

Van Dijck, José (2014). Datafication, dataism and dataveillance: Big Data between scientific paradigm and ideology. *Surveillance & Society*, vol. 12, no. 2, 2014, pp. 197–208.

Vogel, Kathleen M.; Brian L. Balmer; Sam Weiss Evans; Inga Kroener; Miwao Matsumoto; and Brian Rappert (2016). Knowledge and security. In *The Handbook of Science and Technology Studies*. MIT Press.

Wagner, Ben (2021). High Risk AI 9: "AI systems intended to be used for predicting the occurrence of crimes or events of social unrest with a view to allocate resources devoted to the patrolling and surveillance of the territory" (p.68) 23/. *@benwagne_r*. Tweet. https://twitter.com/benwagne_r/status/1382220691127345154. Accessed 31 May 2021

Wang, Shiguang; Prasanna Giridhar; Hongwei Wang; Lance Kaplan; Tien Pham; Aylin Yener; and Tarek Abdelzaher (2017). StoryLine: Unsupervised Geo-event Demultiplexing in Social Spaces without Location Information. In *IoTDI '17. Proceedings of the Second International Conference on Internet-of-Things Design and Implementation*. Pittsburgh, PA, USA: ACM Press, pp. 83–93.

Warner, Michael (2002). Publics and counterpublics. *Public culture*, vol. 14, no. 1, 2002, pp. 49–90.

Weber, Jutta; and Katrin M. Kämpf (2020). Technosecurity Cultures: Introduction. *Science as Culture*, vol. 29, no. 1, January 2020, pp. 1–10.

Williams, Matthew L.; Pete Burnap; and Luke Sloan (2017). Crime Sensing With Big Data: The Affordances and Limitations of Using Open-source Communications to Estimate Crime Patterns. *The British Journal of Criminology*, vol. 57, no. 2, March 2017, pp. 320–340.

Williams, Matthew L.; Adam Edwards; William Housley; Peter Burnap; Omer Rana; Nick Avis; et al (2013). Policing cyber-neighbourhoods: tension monitoring and social media networks. *Policing and Society*, vol. 23, no. 4, December 2013, pp. 461–481.

Williams, Michael J (2008). (In) Security studies, reflexive modernization and the risk society. *Cooperation and Conflict*, vol. 43, no. 1, 2008, pp. 57–79.

Williamson, Kirsty; Lisa M. Given; and Paul Scifleet (2018). Chapter 19 - Qualitative data analysis. In K. Williamson and G. Johanson (Eds.), *Research Methods (Second Edition)*. Chandos Publishing.

Wilson, Steven Lloyd (2017). Detecting Mass Protest through Social Media. *The Journal of Social Media in Society*, vol. Vol 6, No 2, 2017, pp. 21.

Winner, Langdon (1980). Do artifacts have politics? *Daedalus*, 1980, pp. 121–136.

Wu, Congyu; and Matthew S Gerber (2017). Forecasting Civil Unrest Using Social Media and Protest Participation Theory. *IEEE Transactions on Computational Social Systems*, vol. 5, no. 1, 2017, pp. 82–94.

Xu, Jiejun; Tsai-Ching Lu; Ryan Compton; and David Allen (2014). Civil Unrest Prediction: A Tumblr-Based Exploration. In W. G. Kennedy, N. Agarwal, and S. J. Yang (Eds.), *Lecture Notes in Computer Science. Social Computing, Behavioral-Cultural Modeling and Prediction*. Cham: Springer International Publishing, pp. 403–411.

Zhao, Liang; Feng Chen; Chang-Tien Lu; and Naren Ramakrishnan (2015). Dynamic theme tracking in Twitter. In *2015 IEEE International Conference on Big Data (Big Data)*. IEEE, pp. 561–570.

Zhao, Liang; Feng Chen; Chang-Tien Lu; and Naren Ramakrishnan (2016). Online Spatial Event Forecasting in Microblogs. *ACM Transactions on Spatial Algorithms and Systems*, vol. 2, no. 4, November 2016, pp. 1–39.

Zhao, Liang; Qian Sun; Jieping Ye; Feng Chen; Chang-Tien Lu; and Naren Ramakrishnan (2017). Feature Constrained Multi-Task Learning Models for Spatiotemporal Event Forecasting. *IEEE Transactions on Knowledge and Data Engineering*, vol. 29, no. 5, 2017, pp. 1059–1072.

Zhou, Yuchen; Mark Dredze; David A Broniatowski; and William D Adler (2018). Gab: The Alt-Right Social Media Platform, 2018, pp. 6.

Zuboff, Shoshana (2019). *The age of surveillance capitalism: the fight for a human future at the new frontier of power*. New York: PublicAffairs.