Use of Statistics in Disaster by Local Individuals: An Examination of Tweets during COVID-19

Lise Ann St. Denis
CIRES, Earth Lab
University of Colorado, Boulder
Lise.St.Denis@Colorado.edu

Amanda Lee Hughes
Crisis Informatics Lab
Brigham Young University
amanda_hughes@byu.edu

ABSTRACT

We report on how individuals local to the US state of Colorado used statistics in tweets to make sense of the early stages of the COVID-19 pandemic. Tweets provided insight into how people interpreted statistical data, sometimes incorrectly, which has implications for crisis responders tasked with understanding public perceptions and providing accurate information. With widespread concerns about the accuracy and quality of online information, we show how monitoring public reactions to and uses of statistics on social media is important for improving crisis communication. Findings suggest that statistics can be a powerful tool for making sense of a crisis and coping with the stress and uncertainty of a global, rapidly evolving event like the COVID-19 pandemic. We conclude with broader implications for how crisis responders might improve their communications around statistics to the public, and suggestions for how this research might be expanded to look at other types of disasters.

Keywords

Social media, statistics, COVID-19, pandemic.

INTRODUCTION

The COVID-19 pandemic brought about unprecedented uncertainty, as governments and communities worldwide enforced lockdowns to contain the spread of the virus. Social media played a crucial role during this time, serving as a communication channel where people shared and sought information, and expressed their concerns and confusion about the evolving situation (Pine et al. 2021). Examining the messages and conversations on social media platforms can offer valuable insights into how people processed and reacted to the crisis.

In this paper, we present our preliminary findings on the Twitter messages sent by local individuals in Colorado, USA, during the early stages of the COVID-19 stay-at-home orders. Local individuals are often best positioned to provide on-the-ground insights into a crisis event (Bruns and Burgess 2012; Kogan et al. 2015; Starbird et al. 2010), but their voices can get lost in the overwhelming number of social media messages generated during a high-profile and pervasive event like the pandemic (Cinelli et al. 2020). To address this, we used a machine learning classifier previously developed in crisis research (Diaz et al. 2020) to filter out messages from non-local individuals, organizations, businesses, and news media. This resulted in a dataset of 4,352 tweets from local individuals over a two-week period, filtered down from a larger dataset of 101,257 tweets, which is far more manageable for crisis responders to examine. We then manually categorized and analyzed these tweets to better understand how local individuals responded to the crisis and what information could be useful to responders. A primary aim was to identity significant categories within the local individual data that could potentially be extracted using machine learning techniques for future events.

Our analysis revealed a prominent category of tweets that used statistics as a way of making sense of the pandemic and the associated lockdowns. Tweets in this category provided insight into how people interpreted and used statistical data, sometimes incorrectly, which has implications for crisis responders and public officials tasked with understanding public perceptions and providing accurate information. Additionally, we observe that messages containing statistics have semantic properties (i.e., the use of numbers and percentages) or links to specific statistical reporting sites which could make these types of messages easier to identify using machine learning. This presents an opportunity to more easily extract and analyze such information for insights into public perceptions and the potential misuse of data during crisis events. In this paper, we analyze these tweets to identify the uses of statistics by local individuals and the potential for such tweets to improve crisis communication.

BACKGROUND

In the field of crisis informatics, researchers have examined how people communicate about crisis events on social media platforms for 15+ years (Palen and Hughes 2018; Reuter et al. 2018). Members of the public have used social media to find and share critical information during hurricanes (Bica et al. 2019; Hughes et al. 2014), wildfires (Chauhan and Hughes 2017; Waqas and Imran 2019), and floods (Murthy and Longwell 2013). Often the crisis-affected public can provide a unique on-the-ground perspective of the event through their reports and shared videos and photos that can contribute to situational awareness for crisis responders (Vieweg et al. 2010). They also use social media to seek support (e.g., physical support, emotional support) from the local community, as well as from the global audience that social media can provide (Bruns and Burgess 2012; Hughes 2019; Kogan et al. 2015; Palen and Hughes 2018; Reuter et al. 2018).

However, large-scale, high-intensity disasters events often result in millions of social media posts during the height of a disaster's impact, making it impossible to keep pace with the volume of content surrounding an event (Castillo 2016; Imran et al. 2015; Palen and Hughes 2018). To address this, researchers have used machine learning and natural language processing techniques to filter relevant information from the noise (Castillo 2016; Imran et al. 2015; Johnson et al. 2020; Pedrood and Purohit 2018; Purohit et al. 2018). Tools and methods for monitoring social media for emergency response are still limited though because it is difficult to create solutions that work for all hazards and circumstances (Hiltz et al. 2020; Hughes and Shah 2016; Reuter et al. 2018). However, if done well, tools and methods for extracting useful information from social media have the potential to save lives and protect property (Hiltz et al. 2020; Vieweg et al. 2010).

Social Media Data from Local Individuals

A useful approach to filtering the overwhelming amount of data during large-scale events like the COVID-19 Pandemic is to focus on information provided by people local to the event (St. Denis 2015). Local individuals (as opposed to media, advertisers, those far from the event, etc.) tend to provide information that can offer a clearer picture of what is happening at the location of the disaster event and the effects it is having on the affected population (Bruns and Burgess 2012; Starbird et al. 2010). As a result, several studies have focused on extracting local data from social media (Kogan et al. 2015; St. Denis et al. 2020; St. Denis 2015; Starbird and Palen 2012). This study contributes to and builds upon this work by examining how people local to a disaster use statistics to make sense of a disaster event and its impact on their community.

Data and Statistics in Disaster

During a disaster event, people often turn to statistics to help them make sense of the situation (Castillo 2016; Cinelli et al. 2020). Statistics can provide a way to measure and compare the severity of the disaster, the effectiveness of emergency response efforts, and the progress being made towards recovery. Social media platforms have become an important source of this information during disasters, as users often share data and statistics to help others understand the situation (Bica et al. 2019; Chauhan and Hughes 2017; Pine et al. 2021; Starbird et al. 2010). However, there is currently a lack of research that specifically looks at how statistics are used in social media messaging around disaster events. With widespread concerns about the accuracy and quality of information online generally (Del Vicario et al. 2016) and during times of disaster (Chauhan and Hughes 2020; Enders et al. 2020), monitoring how the public reacts to and uses statistics on social media is important for improving crisis communication strategies.

METHOD

In this study, we collected Twitter data from local individuals during the initial COVID-19 response in Colorado using a neural network classifier– the Earth Lab Global Social Sensor (EL GSS). This classifier takes a stream of Twitter data and identifies tweets from the user accounts most likely to contribute new, personalized content during high volume events (Diaz et al. 2020; St. Denis et al. 2020). It does this by using information from the user's account and their recent tweet activity to predict what type of informational role the user plays. Possible informational roles are broken into two types: *Formal* and *General Public*. Formal account roles include the *Media* for mainstream media or media personnel, *Emergency Response* for response organizations/first responders, and *Public Sector* for official city and community organizations and personnel. The *General Public* account roles are divided based on their content. *Personalized* accounts are those likely to contribute first-hand information related to the event, while *Redistribution* accounts primarily synthesize or 'redistribute' information from existing sources. Based on past performance, the user account level classifier is 90% accurate at identifying *Personalized* accounts (St. Denis et al. 2020). We use tweets from *Personalized* accounts in this study because we found that approximately 30% of tweets from *Personalized* accounts contained some form of first-hand information as compared to 6% for accounts labeled as *Redistribution* (St. Denis et al. 2020). We extracted tweets sent from accounts that our classifier labeled as *Personalized* to comprise the data set from local individuals for this study.

We began collecting tweets on April 1, 2020, when one of the researchers on this project was asked by government officials to assist in understanding the public's reaction to the stay-at-home order issued by the governor of Colorado on March 26, 2020 (Polis 2020). At the time, Colorado had "hundreds of confirmed and presumptive cases of COVID-19 and related deaths" and there was also "substantial evidence of community spread of COVID-19 throughout the state" (Ryan 2020 p. 1). The initial stay-at-home order was scheduled to end on April 11, 2020 (Polis 2020), but due to worsening conditions it was extended.

The following terms were used to collect tweets using the Twitter Streaming API: pandemic related terms (covid-19, coronavirus, pandemic) x regional terms (denver, boulder, colorado). The data collection period lasted two weeks—April 1 to April 15, 2020—during which we collected 101,257 tweets. This timeframe captured the initial reaction to the stay-at-home order and provided a snapshot of the public's response to the pandemic.

To analyze the tweets of local individuals, we used the classifier described above to filter down the larger dataset, leaving a subset of 4,352 tweets from individuals who were local to the state of Colorado. With this filtered dataset, we conducted a deeper analysis to identify common themes and patterns.

Twitter Data Analysis

To analyze the local tweets, we employed multiple coders and conducted iterative coding to refine the categories in our codebook. Some of the categories included tweets that discussed impacts on social networks, the availability of testing and personal protective equipment (PPE), and messages about social distancing and its effects. We do not report on all these categories here as we are planning a full reporting for a future publication.

In this work-in-progress paper, we examine a specific category of tweets from our dataset that shed light on the ways in which individuals use statistics to understand the COVID-19 pandemic. Our preliminary analysis revealed interesting behaviors, including using statistics to assess the severity of the pandemic, evaluate the effectiveness of government policies, and understand the risks associated with different activities. We also found a variety of ways in which individuals reported and interpreted statistics. Given the importance of statistics in shaping public perceptions and responses to the pandemic and the ease with which one can "lie" with statistics (Huff 1993), we unpack these sensemaking behaviors.

Table 1: Statistics Tweet Categories and Descriptions

Category	Description
Trend Interpretation	Interprets trends in covid cases, hospitalizations, hospital capacity, and deaths
Regional Comparison	Shares geographic statistics and makes regional comparisons based on these statistics
Data Sharing	Only distributes statistics, no interpretation or reaction to the data
Other Statistics	Discusses specialized populations, antibody testing, economic impacts, and politics

The Statistics Dataset consists of 295 tweets. Any tweet in our local individual dataset that mentioned or described some sort of numerical statistic associated with the Pandemic was included in this set. Two coders reviewed each of the tweets in the dataset and iteratively refined a coding scheme for this data (see Table 1). Since the dataset was small, both coders reviewed all tweets and discussed any discrepancies between codes until agreement could be reached.

RESULTS

Figure 1 reports the number of tweets found in each of the statistical discussion categories. The largest category was *Trend Interpretation* (106). Tweets in this category discussed general trends in cases, hospitalizations, and deaths. We also observed a significant number of tweets (74) that examined trends in case data relative to other regions of the country or local areas. These tweets fell in the *Regional Comparison* category. Some tweets simply engaged in *Data Sharing* (63) where they distributed current information sources or statuses without interpretation or analysis. Finally, the *Other Statistics* category contains tweets (52) that focused on unique populations or interests. The two primary reference sites that we found across all statistical tweets were the Colorado Department of Public Health and Environment's Covid-19 Reporting Page (https://covid19.colorado.gov/case-data) and the Covid-19 projections provided by the Institute for Health Metrics and Evaluation (IHME) (https://covid19.healthdata.org). We now describe each of these statistical tweet categories in more detail.

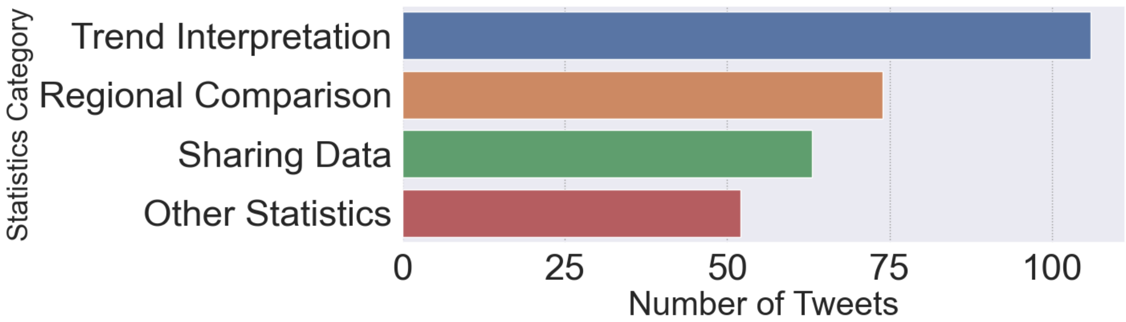

Figure 1: Number of Tweets in Each Statistics Category

Trend Interpretation

In the initial weeks of the stay-at-home orders, information about the novel virus was scarce with individuals interpreting updates in social isolation with limited information. The models that projections were based on were subject to adjustment as the science evolved in the early weeks of the pandemic. Some commentary in our dataset discussed the variability across these models and the confusion surrounding how to interpret the updated projections.

When analyzing the *Trend Interpretation* tweets, we noticed a significant degree of polarity in the messaging. To gain a better understanding of this phenomenon, we conducted further coding of these tweets, classifying them as either positive, negative, or neutral in their interpretation based on the text of the tweet itself (see Figure 2). For example, if a tweet mentioned "good news" or "flattening the curve," it was labeled as positive. Conversely, if a tweet referenced "bad news" or a rise in hospitalizations, it was classified as negative. If the sentiment was unclear or the tweet did not express a clear positive or negative sentiment, we labeled it as neutral. Overall, there were more positive tweets (43) than negative (31), and fewer neutral tweets (20). Our analysis revealed a fairly even distribution of positive and negative tweets on average across the reporting period, highlighting the diverse interpretations of COVID-19 statistics by locals. Understanding these varied perspectives can help crisis responders and public health officials develop effective communication strategies that address public concerns and misconceptions.

The sentiment of *Trend Interpretation* tweets fluctuated over time (see Figure 2). For example, on April 3rd and 4th the posts were primarily negative. An example tweet reads:

> *"Have been watching these numbers from here in #Colorado creep closer together. Not every hospitalized COVID patient is necessarily in an ICU bed, but as the numbers get closer, we near overflow. And there are still non-COVID hospitalizations. Scary."*

However, over the next few days, the posts shifted, with many tweets highlighting positive developments. Many of these positive tweets mentioned a new model, which suggested that the peak of fatalities had passed:

> *"Oh Wow. The prior IHME model had the peak in Colorado fatalities in late April, and dire predictions for hospitals. Now, with more data, the expectation value is that we have passed the peak, and the upper 95% confidence interval is April 12 for the peak. <link to the data>"*

The release of more optimistic models and data appeared to be associated with an increase in positive tweets.

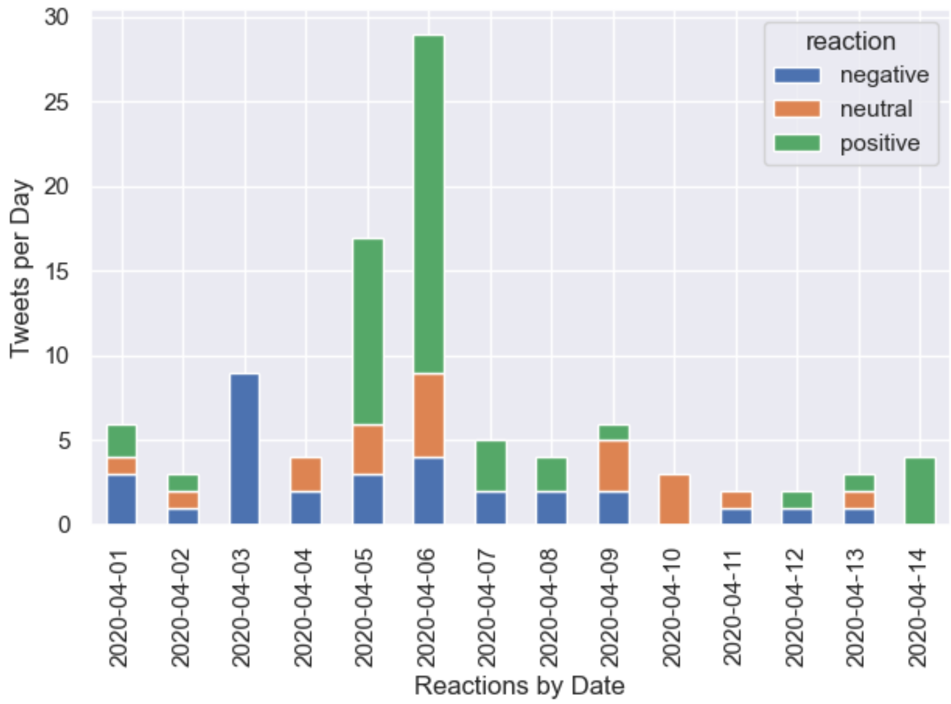

Figure 2: Trend Interpretation tweets broken down by sentiment (negative, positive, Neutral) over time

In *Trend Interpretation* tweets, we observed that individuals found it challenging to draw clear conclusions from daily fluctuations in the raw reporting data provided by state health agencies. For instance, we found two tweets from the same day that interpreted the same information in vastly different ways (refer to Table 2). While one individual saw the information as evidence of "flattening the curve," the other perceived it as a sign of a surging death rate. Some of the tweets we labeled as neutral were conversations about the reasons behind daily anomalies in the data, such as testing/reporting lags that occur over the weekend. Still others reflected the difficulties of trying to make sense of an unprecedented situation that was difficult to read. We didn't know then how difficult interpreting these curves would continue to be.

Table 2: Interpretation of General Case Trends - Example of Positive, Negative, and Neutral Tweets

Positive **Post from April 8th:**

"@<reporter>...There weren't 29 new deaths all in one day, they were updated numbers over several previous days. The curve appears to be flattening - check it out. You have to look at the most recent bar graph on the state site. [url]"

Negative **Post April 8th:**

"Colorado #Coronavirus update for 4/7:

-- 29 deaths, far exceeding our previous high in a day (18 on March 31st)
-- 257 confirmed new cases, up from yesterday's 222, but still only a 4.97% cases growth rate (which is good!)
-- Death rate skyrocketed to 3.30% (exceeding USA)"

Neutral **Discussion Surrounding Reporting Bias on April 4th:**

"Seems to be a bias in the Colorado COVID-19 case-by-onset-date data where fewer people report an onset COVID-19 on the weekends?
I assume this is related to less testing on weekends coupled with test dates being reported as onset dates when onset date is lacking, but 🤷. <link to state reporting site>"

Regional Comparison

Local individuals often compared their region's COVID-19 situation with that of other regions, particularly in the early days of the pandemic when there was much uncertainty. Most comparisons were between Colorado and other states but sometimes comparisons were made between counties and cities within Colorado. Typical tweets in this category showed how Colorado's numbers ranked nationally:

> *"#Coronavirus cases, #USA New York: 113,704 New Jersey: 34,124 Michigan: 14,225 California: 13,878 Louisiana: 12,496 Massachusetts: 11,736 Florida: 11,545 Pennsylvania: 10,444 Illinois: 10,357 Washington: 7,591 Texas: 6,739 Georgia: 6,383 Connecticut: 5,276 Colorado: 4,565"*

> *"@<news organization> As of today, Nationally, and per capita, Colorado ranks #12 in COVID deaths, #14 in confirmed COVID cases, and #38 in people tested."*

These tweets provided a way for people to gauge how well Colorado was managing the pandemic compared to other areas and gain a broader context of the severity of the situation in their local community. By comparing their state or city with others, people could also draw on strategies implemented in other places and assess their effectiveness. This helped to inform discussions around policy decisions, such as the timing and extent of reopening measures, and provided a way for people to express their opinions on these matters.

In addition, regional comparisons provided a way for people to hold leaders accountable and advocate for necessary actions. By highlighting how their region compared to others, people could call attention to disparities and push for more resources and support from state and federal governments. This was particularly important for areas where the virus spiked early and signaled significant concern. For example, Eagle County, Colorado was compared to global and the United States hotspots in terms of COVID-19 cases due to its early number of positive cases and exposure risk attributed to foreign tourism traffic:

> *"4/2/20 #coronavirus numbers of confirmed cases per capita:*
> *World: 1 of every 7,570 people*
> *U.S.: 1 of every 1,353 people*
> *New York state: 1 of every 212*
> *Colorado: 1 of every 1,528*
> *Eagle County, Colorado: 1 of every 174"*

The data shows that Eagle County had a higher number of COVID-19 cases per capita than the US state of New York, which is surprising given that New York was one of the earliest and hardest-hit areas during the pandemic. This underscores the need to consider factors beyond surface-level data, such as population density, social context, and testing rates, when making comparisons between regions.

Some local individuals did not fully grasp the importance of normalizing numbers by population for accurate comparisons. For instance, one tweet author expressed confusion about the practice of adjusting COVID-19 death numbers by population:

> *"@<user #1> @<user #2> @<user #3> Indiana has a population of 6.7 million. Maryland 6 mil. Ohio, 11 m. Colorado 5.8 m. Texas, 29 m. And each one of them individually have over 300 deaths. None of those states are even in the top 10 of US COVID deaths. Tell me again what total population has to do with anything?"*

Responses like these indicate a lack of statistical literacy among the public, which can lead to misinformation and misunderstandings about the severity of the pandemic. This lack of understanding can be particularly challenging for crisis responders during a disaster, as they must ensure that the public has a clear and accurate understanding of the situation. It is therefore essential for responders to account for this lack of understanding in their messaging and communication strategies. One way to do this is by providing clear and concise explanations of statistical concepts, such as normalization by population, and why they are important for accurate comparisons. This can help to combat misinformation and promote accurate understanding among the public.

Data Sharing

Tweets that distributed statistics only, without interpretation were labeled as Data Sharing. We identified 63 tweets in this category. Many of these linked to the two primary statistics websites above and/or shared the latest case statistics. The following tweet is an example of a daily update with a link to the state website:

> *"Updated Colorado Numbers (Updated 4-1, data from end of day 3-31) reported 3/31 →4/1*
> *Cases 2966→3342*
> *Hospitalized: 509→620*
> *Counties: 50→50 (out of 64)*
> *Deaths 69–80*
> *covid19.colorado.gov/case-data"*

The main objective of individuals that posted tweets in this category appeared to be the dissemination of information and the expansion of the reach of COVID-19 data among the poster's social network. However, it is important for crisis responders to monitor these messages for accuracy because they can spread widely if not corrected, contributing to the spread of misinformation.

Other Statistics

We found a small sample of tweets focused on other statistical discussions. This included discussion of outbreaks and risks to particular populations including prison populations, essential workers, nursing homes, the elderly, and risks based on ethnicity or gender. There was an outbreak in a local meatpacking plant affecting over 300 workers that factored prominently in the news during the study period. Another media story discussed the low incidence of COVID-19 detected in a rural Colorado county (less than 1%) based on widespread antibody testing provided by a local manufacturer. While some interpreted this as encouraging news, others point out testing results are unlikely to be representative of more populous areas:

> *"Here is the website with the raw #coronavirus nasal swab testing data vs antibody testing data. Important to note that it comes from San Miguel County, Colorado which may not reflect the prevalence in a "Hot Spot" like NY City <link>"*

A few tweets in the *Other Statistics* category made a direct connection to the possible economic impacts of the shutdown in Colorado (e.g., loss of wages, inability to pay rent), but at the time it was too early to assess these impacts. There were also tweets that made a direct connection between case data and politics. In one Northern Colorado county where cases were rising, constituents were critical of the slow response from political representatives:

> *"Latest data from @CDPHE: Weld County, population 324,492, is #1 for #COVID19 deaths. More than Denver, population 693,417. @RepKenBuck of Weld County is still scoffing at protective measures--today. <link #1> How can this not outrage every one of us? #copolitics <link #2>"*

Lastly in this category, several tweets contained references to known misinformation campaigns such as #filmyourhospital. The #filmyourhospital movement was a social media campaign that emerged in the early days of the COVID-19 pandemic (Ahmed et al. 2020). The movement encouraged people to visit their local hospitals and record videos or take photos to prove that the hospitals were not as overwhelmed as the media and government officials were claiming. Unfortunately, statistics can be used to propagate misinformation like this, leading to distrust in government and health officials. Monitoring social media for the spread of misinformation can help these officials understand what messages are circulating in the public sphere and take action to correct them.

DISCUSSION & FUTURE WORK

In this study, we observed that people used statistics in different ways to make sense of the COVID-19 pandemic. Some individuals used statistics to evaluate the severity of the pandemic and how to respond, while others used statistics to challenge official reports or to criticize government actions. Additionally, we found that people used statistics to express emotions such as fear, frustration, and hope. Our findings suggest that statistics can serve as a tool for individuals to make sense of a crisis and to cope with the associated stress and uncertainty of a global, rapidly evolving event like the COVID-19 pandemic.

Our analysis also reveals that people interpretted statistics in diverse ways during crises. While some viewed the declining number of COVID-19 cases positively, suggesting that the pandemic was being brought under control, others viewed the same statistics negatively, indicating that the pandemic was raging out of control. These divergent interpretations emphasize the importance of crisis responders monitoring social media to understand public perceptions and communicating statistics effectively. By understanding the various ways that people interpret statistics and tailoring their messaging accordingly, crisis responders can ensure that accurate information is disseminated and that the public is empowered to make informed decisions.

The use of statistics is not unique to the COVID-19 pandemic, and comparisons can be drawn to statistics use during other types of disaster events. For instance, during natural disasters like hurricanes people often rely on statistics such as the number of evacuations, the amount of damage, and the likelihood of the projected path of the storm to make decisions about how to prepare and respond. Wildfires have similar types of statistics around evacuations and damages, but also have statistics about the number of acres burned or the percentage of containment. During mass shootings, people may use statistics to understand the prevalence of gun violence and to advocate for policy change. Analyzing the types of statistics shared and used across social media in different disaster events could allow us to create a taxonomy across event types and better understand their unique features, which could in turn help in developing guidelines and materials for crisis responders to communicate more effectively during these events. A unique aspect of our study was the use of a machine learning classifier to identify and analyze local communication patterns around the COVID-19 pandemic, and this approach could be extended to other regions and crisis situations to better understand how people use statistics in different contexts. By understanding which statistics are most meaningful to the public and the potential risks to data quality and misinterpretation, managers can tailor their messaging and better meet the information needs of those affected by disaster.

The focus of this paper has been on the use of statistics in social media. While statistics have a high potential for easy extraction using machine learning, we recognize that there are likely other types of informational categories shared by local individuals that could be useful for crisis responders to understand and strategize for. For example, expressions of mistrust in the government or crisis response officials could be important indicators of potential challenges in the response efforts. In future studies, we hope to explore other types of information shared local individuals.

The EL GSS classifier has continued to evolve since we began work on this study. We can now both identify the informational role and type of content being shared using a similar categorization to the informational role. This allows us to focus on localized discussions around official sources, filtering out the redistribution of media coverage and newsfeeds (St. Denis 2022). We are planning a second round of analysis focused on conversational threads related to tweets from emergency response and public sector accounts. Understanding how the public reacts to and contribute to content shared by officials will complement our current work. We are also considering how this information could be used to identify and monitor for misinformation and polarizing responses. To do this we anticipate extending the classifier to include deep learning topic models (Müller et al. 2020; Xu et al. 2022).

Finally, we are working on a live-stream prototype that will filter out much of the noise and provide crisis responders or public officials with locally rich content. We plan to evaluate this prototype through a series of user studies, which will help us further refine and evaluate how localized public response could be used to shape ongoing updates during a future pandemic or complex event. Overall, our goal is to provide more accurate and effective communication during crises to help the public make informed decisions and reduce the impact of disasters.

Limitations

We note that the sample size of tweets used in this study was small, which could limit the generalizability of our findings. Future studies could collect data from a larger sample size. Additionally, our study only examined the use of statistics during the initial stages of the COVID-19 pandemic, and this behavior likely changed over time as cases rose and fell and vaccines became available. Therefore, it would be worthwhile to investigate how people's reactions to statistics evolved as the pandemic progressed. Furthermore, we recognize that our focus on Twitter is a limitation of this study. The classifier we used was developed using Twitter data because this data is publically and freely available through an accessible API, however it is important to recognize that social media is much broader than Twitter. While other platforms may not be as accessible, they may provide better sources of local information based on how they are used, the users they attract, and the types of communities they support. We anticipate that the techniques used in our classifier to identify local individuals could be transferable to other social media platforms, and we plan to explore this in future work.

ACKNOWLEDGEMENTS

Funding was provided by Earth Lab through the University of Colorado Boulder's Grand Challenge Initiative and the Cooperative Institute for Research in Environmental Science (CIRES).

REFERENCES

Ahmed, W., F. L. Seguí, J. Vidal-Alaball, and M. S. Katz. 2020. "COVID-19 and the 'Film Your Hospital' Conspiracy Theory: Social Network Analysis of Twitter Data." *Journal of Medical Internet Research*, 22 (10): e22374. https://doi.org/10.2196/22374.

Bica, M., J. L. Demuth, J. E. Dykes, and L. Palen. 2019. "Communicating Hurricane Risks: Multi-Method Examination of Risk Imagery Diffusion." *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 1–13. New York, NY: ACM.

Bruns, A., and J. E. Burgess. 2012. "Local and Global Responses to Disaster: #eqnz and the Christchurch Earthquake." *Proceedings of the Disaster and Emergency Management Conference*, 86–103. Brisbane, QLD: AST Management Pty Ltd.

Castillo, C. 2016. *Big Crisis Data: Social Media in Disasters and Time-Critical Situations*. New York, NY, USA: Cambridge University Press.

Chauhan, A., and A. L. Hughes. 2017. "Providing Online Crisis Information: An Analysis of Official Sources during the 2014 Carlton Complex Wildfire." *Proceedings of the 2017 International Conference on Human Factors in Computing Systems*, 3151–3162. New York, NY: ACM.

Chauhan, A., and A. L. Hughes. 2020. "Trustworthiness Perceptions of Social Media Resources Named after a Crisis Event." *Proceedings of the ACM on Human-Computer Interaction*, 4 (CSCW1): 1–23. https://doi.org/10.1145/3392849.

Cinelli, M., W. Quattrociocchi, A. Galeazzi, C. M. Valensise, E. Brugnoli, A. L. Schmidt, P. Zola, F. Zollo, and A. Scala. 2020. "The COVID-19 social media infodemic." *Scientific Reports*, 10 (1): 16598. Nature Publishing Group. https://doi.org/10.1038/s41598-020-73510-5.

Del Vicario, M., A. Bessi, F. Zollo, F. Petroni, A. Scala, G. Caldarelli, H. E. Stanley, and W. Quattrociocchi. 2016. "The Spreading of Misinformation Online." *Proceedings of the National Academy of Sciences*, 554–559.

Diaz, J., L. A. St. Denis, M. B. Joseph, K. Solvik, and J. K. Balch. 2020. "Classifying Twitter Users for Disaster Response: A Highly Multimodal or Simple Approach?" *Proceedings of the Information Systems for Crisis Response and Management Conference (ISCRAM 2020)*.

Enders, A. M., J. E. Uscinski, C. Klofstad, and J. Stoler. 2020. "The Different Forms of COVID-19 Misinformation and Their Consequences." *The Harvard Kennedy School (HKS) Misinformation Review*, 1 (8). https://doi.org/10.37016/mr-2020-48.

Hiltz, S. R., A. L. Hughes, M. Imran, L. Plotnick, R. Power, and M. Turoff. 2020. "Exploring the Usefulness and Feasibility of Software Requirements for Social Media Use in Emergency Management." *International Journal of Disaster Risk Reduction*, 42: 101367. https://doi.org/10.1016/j.ijdrr.2019.101367.

Huff, D. 1993. *How to Lie with Statistics*. W. W. Norton & Company.

Hughes, A. L. 2019. "Site-Seeing in Disaster: Revisiting Online Social Convergence a Decade Later." *New Media in Times of Crisis*, K. K. Stephens, ed., 193–208. New York, NY: Routledge.

Hughes, A. L., and R. Shah. 2016. "Designing an Application for Social Media Needs in Emergency Public Information Work." *Proceedings of the 19th International Conference on Supporting Group Work*, GROUP '16, 399–408. New York, NY, USA: ACM.

Hughes, A. L., L. A. St. Denis, L. Palen, and K. M. Anderson. 2014. "Online Public Communications by Police & Fire Services during the 2012 Hurricane Sandy." *Proceedings of the 2014 International Conference on Human Factors in Computing Systems*, 1505–1514. New York, NY: ACM.

Imran, M., C. Castillo, F. Diaz, and S. Vieweg. 2015. "Processing Social Media Messages in Mass Emergency: A Survey." *ACM Comput. Surv.*, 47 (4): 67:1-67:38. https://doi.org/10.1145/2771588.

Johnson, M., D. Murthy, B. Roberstson, R. Smith, and K. Stephens. 2020. "DisasterNet: Evaluating the Performance of Transfer Learning to Classify Hurricane-Related Images Posted on Twitter."

Kogan, M., L. Palen, and K. M. Anderson. 2015. "Think Local, Retweet Global: Retweeting by the Geographically-Vulnerable During Hurricane Sandy." *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, 981–993. New York, NY: ACM.

Müller, M., M. Salathé, and P. E. Kummervold. 2020. "Covid-twitter-bert: A Natural Language Processing Model to Analyse Covid-19 Content on Twitter." *arXiv preprint arXiv:2005.07503*.

Murthy, D., and S. A. Longwell. 2013. "Twitter and Disasters: The Uses of Twitter during the 2010 Pakistan Floods." *Information, Communication & Society*, 16 (6): 837–855.

https://doi.org/10.1080/1369118X.2012.696123.

Palen, L., and A. L. Hughes. 2018. "Social Media in Disaster Communication." *Handbook of Disaster Research*, Handbooks of Sociology and Social Research, H. Rodriguez, W. Donner, and J. E. Trainor, eds., 497–518. Springer International Publishing.

Pedrood, B., and H. Purohit. 2018. "Mining Help Intent on Twitter During Disasters via Transfer Learning with Sparse Coding." *Social, Cultural, and Behavioral Modeling*, Lecture Notes in Computer Science, R. Thomson, C. Dancy, A. Hyder, and H. Bisgin, eds., 141–153. Cham: Springer International Publishing.

Pine, K. H., M. Lee, S. A. Whitman, Y. Chen, and K. Henne. 2021. "Making Sense of Risk Information amidst Uncertainty: Individuals' Perceived Risks Associated with the COVID-19 Pandemic." *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 1–15. New York, NY: ACM.

Polis, J. 2020. "Governor Polis Announces Statewide Stay-At-Home Order, Provides Update on Colorado Response to COVID-19." https://www.colorado.gov/governor/news/gov-polis-announces-statewide-stay-home-order-provides-update-colorado-response-covid-19.

Purohit, H., C. Castillo, M. Imran, and R. Pandev. 2018. "Social-EOC: Serviceability Model to Rank Social Media Requests for Emergency Operation Centers." *2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, 119–126. IEEE.

Reuter, C., A. L. Hughes, and M.-A. Kaufhold. 2018. "Social Media in Crisis Management: An Evaluation and Analysis of Crisis Informatics Research." *International Journal of Human–Computer Interaction*, 34 (4): 280–294. https://doi.org/10.1080/10447318.2018.1427832.

Ryan, J. H. 2020. *Implementing Stay at Home Requirements*. Amended Public Health Order. Colorado Department of Public Health & Environment.

St. Denis, L. A. 2015. "Emergency Responders as Inventors: An Action Research Examination of Public Information Work." PhD Dissertation. University of Colorado at Boulder.

St. Denis, L. A. 2022. "Social Media Content Filtering For Emergency Management." U.S. Patent Application No: 18/000,314. Filed: 11/30/2022.

St. Denis, L., A. Hughes, J. Diaz, K. Solvik, M. Joseph, and J. K. Balch. 2020. "What I Need to Know is What I Don't Know!": Filtering Disaster Twitter Data for Information from Local Individuals." *Proceedings of the 17th International Conference on Information Systems for Crisis Response and Management, Blacksburg, VA, USA*, 24–27.

Starbird, K., and L. Palen. 2012. "(How) Will the Revolution be Retweeted?: Information Propagation in the 2011 Egyptian Uprising." *Proceedings of the 2012 Conference on Computer Supported Cooperative Work*, 7–16. New York, NY: ACM.

Starbird, K., L. Palen, A. L. Hughes, and S. Vieweg. 2010. "Chatter on the Red: What Hazards Threat Reveals about the Social Life of Microblogged Information." *Proceedings of the ACM 2010 Conference on Computer Supported Cooperative Work (CSCW 2010)*, 241–250. New York, NY: ACM.

Vieweg, S., A. L. Hughes, K. Starbird, and L. Palen. 2010. "Microblogging During Two Natural Hazards Events: What Twitter May Contribute to Situational Awareness." *Proceedings of the ACM 2010 Conference on Computer Human Interaction*, 1079–1088. New York, NY: ACM Press.

Waqas, H., and M. Imran. 2019. "#CampFireMissing: An Analysis of Tweets About Missing and Found People From California Wildfires." *Proceedings of the 16th International Conference on Information Systems for Crisis Response And Management*.

Xu, W. W., J. M. Tshimula, È. Dubé, J. E. Graham, D. Greyson, N. E. MacDonald, and S. B. Meyer. 2022. "Unmasking the Twitter Discourses on Masks During the COVID-19 Pandemic: User Cluster–Based BERT Topic Modeling Approach." *JMIR Infodemiology*, 2 (2): e41198. JMIR Publications Toronto, Canada. https://doi.org/10.2196/41198.