



Pulling through together: social media response trajectories in disaster-stricken communities

Danaja Maldeniya¹  · Munmun De Choudhury⁴ · David Garcia^{5,6,7} · Daniel M. Romero^{1,2,3}

Received: 30 March 2022 / Accepted: 1 April 2023
© The Author(s) 2023

Abstract

Disasters are extraordinary shocks that disrupt every aspect of the community life. Lives are lost, infrastructure is destroyed, the social fabric is torn apart, and people are left with physical and psychological trauma. In the aftermath of a disaster, communities begin the collective process of healing, grieving losses, repairing damage, and adapting to a new reality. Previous work has suggested the existence of a series of prototypical stages through which such community responses evolve. As social media have become more widely used, affected communities have increasingly adopted them to express, navigate, and build their response due to the greater visibility and speed of interaction that these platforms afford. In this study, we ask if the behavior of disaster-struck communities on social media follows prototypical patterns and what relationship, if any, these patterns may have with those established for offline behavior in previous work. Building on theoretical models of disaster response, we investigate whether, in the short term, community responses on social media in the aftermath of disasters follow a prototypical trajectory. We conduct our analysis using computational methods to model over 200 disaster-stricken U.S. communities. Community responses are measured in a range of domains, including psychological, social, and sense-making, and as multidimensional time series derived from the linguistic markers in tweets from those communities. We find that community responses on Twitter demonstrate similar response patterns across numerous social, aspirational, and physical dynamics. Additionally, through cluster analysis, we demonstrate that a minority of communities are characterized by more intense and enduring emotional coping strategies and sense-making. In this investigation of the relationship between community response and intrinsic properties of disasters, we reveal that the severity of the impact makes the deviant trajectory more likely, while the type and duration of a disaster are not associated with it.

Keywords Disasters · Recovery trajectories · Social media · Time series

Extended author information available on the last page of the article

Introduction

Disasters have a multitude of effects on communities, including psychological and physical trauma, as well as disruption of social structures [1, 2]. Research has shown that communities can respond in a variety of different ways to these exogenous shocks [3]. Some prove to be resilient and are able to bounce back to their pre-disaster normal or even become even more well adjusted and tightly knit. For other communities, the effects of a disaster can be devastating, inducing negative effects on individual, social, and economic fronts [1]. Additionally, the outcome for a community after a disaster is often mediated by the nature of the disaster itself: was there forewarning? Do similar events occur frequently in the region? How severe was the initial impact on the community? Was it a prolonged hazard, such as a wildfire, or a sudden but short-lived disaster, such as a tornado? Answers to these questions can be the difference between a community that overcomes adversity to thrive or one that disintegrates [2–4].

In exploring these different questions, researchers have thoroughly analyzed how communities respond to disasters in terms of changes in affect, cognitive processing, social dynamics, and other dimensions [1, 2, 5]. In addition, attempts have been made to develop theoretical frameworks that explain how the immediate effects of a disaster and the subsequent response of the community develop over time [6–9]. For a single member of an affected community, the road to recovery can take many forms, including (i) chronic negative psychological, physiological, and social effects that last for years, (ii) negative effects that become more visible over time, or even (ii) a positive response despite adversity [10]. Previous work has discussed how these individual socially embedded reactions to a disaster may shape the response of the larger community. Theories of disaster response, as well as empirical studies of disaster-stricken communities, propose that communities may go through a characteristic series of psychological and social stages as part of a “response trajectory” [6, 7, 9].

Community Response on Social Media Over the past 2 decades, social media have become the dominant medium through which people, particularly in developed countries, express opinions, consume information, and engage socially. In the context of disasters, this has meant widespread use of social media as a tool and a medium to navigate and respond to the aftermath within and outside affected communities [11–14]. Social media have several advantages (and limitations) that are markedly different from the modes of communication and social interaction that preceded them. These advantages have influenced their use in the aftermath of disasters. Key among these, social media allow regular citizens to engage rapidly with large audiences and provide tools for parsing large volumes of messages to identify relevant information (e.g., hashtags, search functionality). As a consequence, social media, such as Twitter and Facebook, represent the primary, but not only, channel through which those outside of a disaster-stricken community—such as the general public, traditional news media, and relief agencies—perceive its status and ongoing response to disaster. It is important to understand to what extent the collective social media behavior of a community

affected by a disaster can provide a holistic view of itself in the aftermath of disaster. Do community responses observed on social media contain regular patterns that generalize across different events? How is that behavior similar to or different from response patterns that were identified from offline behavior within such communities before the widespread adoption of social media? In this regard, the existing empirical work falls short.

Studies that have explored the temporal evolution of community response on social media have limited themselves to a few of them [13, 15, 16]. We are unaware of any research that has explored the existence of patterns that are relevant in a wide range of disasters. Research along this avenue would have a number of implications. First, if it is established that community responses on social media follow regular patterns, it would allow us to better understand and track new disaster-stricken communities in a sign-posted timeline of their response in the aftermath. To the extent that the social media behaviors of affected communities can be mapped to their real experiences, the ability to track the community response can be very useful for organizations involved in disaster relief and recovery operations. Additionally, this would allow future research to investigate the relationship between contextual variables, such as community attributes, media attention, and support, and deviations in the social media response of individual communities from prototypical patterns. Finally, we expect that the extent to which communities talk on social media about different aspects of their experience will reveal which of them are highlighted in the outside or public consciousness, as well as important gaps in that perception.

Research questions We begin with the existing understanding of the response of the community to disasters. A survey of existing work reveals several common themes [6, 7, 9]. In the beginning of a disaster, people focus first on the safety of themselves and their families. Next, they rescue and help others in the community before any external assistance arrives. In turn, these altruistic behaviors lead to a short-lived period of enhanced community identity that suppresses some of the negative psychosocial effects of the disaster. However, eventually, communities face long-term ramifications of the disaster that lead to a resurgence of those effects, and community members begin the long road to establish a new normal. Our first goal in this study is to establish whether community responses to disasters that are observed on social media follow a prototypical trajectory over time.

RQ1(a) Does community response reflected on social media exhibit a prototypical trajectory across a wide range of disasters?

If it does emerge that there are distinct patterns in community responses across a variety of disasters, those observations may or may not be explained by the previous work, which brings us to a related question.

RQ1(b) Do the temporal patterns of the community response to disasters observed on social media align with the existing stage models of disaster response?

Finally, given the numerous intrinsic differences between natural hazards and communities considered in our study, we expect that community responses will show some differences across disasters. Although each disaster is a unique event, previous work has explored the influence of coarse categorizations of disasters, such as type, severity, and duration, on how affected communities recover over time. This leads to our second research question.

RQ2 Are there distinguishable differences in community response trajectories across different disasters and are these correlated with broad disaster categorizations?

The present work In this study, we begin by engaging with the existing literature to understand how individuals and communities respond to disasters and to consider multiple dimensions of their experiences. Based on this survey, we develop a comprehensive framework of five broad themes (psychological, social, sense-making, biological, and physical) as the lens through which we observe community response. We apply this framework to Tweets from more than 200 U.S. communities affected by disasters to quantify their short-term disaster response as multidimensional time series. We represent each community by Twitter users physically present in the affected area during the disaster, which amounts to more than 2 million users and more than 200 million tweets in all disasters. Finally, we cluster the response trajectories of different communities to establish if there are broadly applicable characteristic patterns of disaster response and if any variations upon these patterns are associated with properties of the specific hazard.

Contributions This study makes three main contributions. First, we have developed an approach that characterizes community response to disasters over social media as a series of prototypical trajectories that are more temporally fine-grained and precise compared to theoretical models and empirical studies of individual events that precede it; we are able to unpack prototypical behavior over days instead of weeks or months. Furthermore, with this approach, we provide a more nuanced narrative of community response by simultaneously considering the many dimensions of that process, which is a necessary stepping stone to understanding how different aspects of community experience, such as social and emotional dynamics, interact and contribute to eventual outcomes.

Second, our findings reveal that, despite wide variations in the properties of the hazards and affected communities, many aspects of community response reflected on social media (social, aspirational, biological, and physical) follow broadly prototypical paths over the short term. These observations on social media broadly align with and validate prior theories on disaster response trajectories. Beyond confirming that community behavior over social media displays regular temporal patterns in the aftermath of disasters, our results reveal a number of previously unobserved themes in the community response. First, we observe a marked decline in conversation that implied productivity. This coincides with the sense of risk and uncertainty that comes with disaster. Second, the onset of disaster coincides with people's focus contracting to their immediate social circle, a phenomenon that gradually abates over time, returning social engagement to pre-disaster levels during our study

period. Third, while we observe a sudden decline followed by a gradual recovery of attention toward a number of aspects of life such as *health*, *finances*, *work*, and *food*, we note that references to *food*-related dynamics take substantially longer to recover. As previous work has shown that the state of food systems is a reliable indicator of community health and sustainability, this observation hints at the potential utility of the trajectory of *food*-related attention measured on Twitter to evaluate community recovery.

Finally, we show that not all behaviors considered in trajectory models are universally observed in disasters, a fact supported by observations from previous disaster research that certain behaviors, such as coping strategies, can vary depending on the particulars and nature of a disaster or the affected community [1]. In our case, a minority of our disasters (25%) exhibit a substantially heightened level of emotions. These communities had very high levels of anger, incredulity, anxiety, and fear at the start of the disaster compared to others. Additionally, these communities experience a resurgence of fear and anger within a few weeks after an initial burst at the onset of disaster has abated. This has implications for organizations that provide relief and recovery assistance, as negative affects such as anger and anxiety play a mediating role in long-term recovery outcomes in communities [17].

The resurgence of fear and anger coincides with the continued signs of elevated cognitive processing associated with *sense-making* within these same communities. These observations indicate an ongoing struggle to make sense of and adapt to a *new normal*. The possibility that these communities are struggling in the aftermath of a disaster, more so than the majority of communities in our dataset, is reinforced by the fact that the amount of aid received by these communities, which is often considered a proxy for the destructiveness of a disaster, is on average higher for these communities after accounting for the type of disaster and duration.

Background and related work

Dimensions of disaster response

Previous work has identified a wide range of different dimensions along which communities express their response to disasters [1, 2, 18, 19]. We take advantage of this literature to establish a thematic framework that allows us to build a comprehensive picture of disaster response in a tractable manner. Our framework is organized along five broad themes: (i) psychological, (ii) psychosocial, (iii) sense-making, (iv) biological, and (v) physical.

Psychological

Disasters impose substantial psychological stress on individuals and communities; affected individuals must deal with fear for their safety, loss of loved ones and homes, and disruption of their social networks and support systems. Previous work has theorized and, in some cases, empirically studied a range of different

psychological manifestations in individuals as they navigate these challenging circumstances. Affected individuals may manifest this distress in the form of post-traumatic stress disorder, depression, anxiety, and other related disorders [1, 6, 19–22]. Although previous work has consistently documented the increased prevalence of these disorders in disaster-stricken communities, there is substantial disagreement on the magnitude of this effect, with estimates showing a wide variance [1]. The level of psychological distress experienced within a specific community is likely to be influenced by a variety of variables, such as disaster characteristics, community demographics, and social resources.

Our interest lies in measuring the general psychological well-being of the community rather than in identifying the prevalence of specific psychological disorders. As such, we limit our attention to the affective and cognitive dimensions of expression that are commonly associated with psychological well-being and are observed within disaster-stricken communities. For example, people diagnosed with one or more of a variety of clinically diagnosed psychological disorders, such as PTSD, depression, anxiety, and other mood disorders, are likely to exhibit elevated levels of negative affect, in the form of sadness, anger, or irritability, anxiety, and fear, among others [1, 23]. Diagnostic definitions of these disorders also include cognitive and behavioral symptoms, such as loss of energy and motivation, withdrawal, and helplessness.

Social

Social dynamics within a community before, during, and after a disaster are central to disaster response. In fact, increasingly, interpretations of disasters place emphasis on their socially constructed nature [3, 24, 25]. The social response to disaster, during and after it, can be considered in a number of related dimensions. Previous work has shown that the aftermath of a disaster can bring a community together. People are more likely to highlight their membership in the community and refer to that affiliation positively [6, 9]. They may take advantage of that shared affiliation to exhort other community members to unite to withstand the disaster. In turn, this shared sense of community can increase social cohesion and lead to a more efficient and effective community response. Additionally, in the early stages of disaster recovery, community members are more likely to seek social support, be more prosocial, and develop solidarity. However, in the long term, the ability of communities to provide support may be exhausted, despite many affected people remaining in need [1, 6]. A more narrow and personal aspect of the social dimension of community response is that individual concern for the safety and well-being of immediate family members, friends, and neighbors is a strong driving force behind behavior during the early stages of a disaster [1, 6, 26].

Sense-making

At the beginning, during, and after a disaster, members of the stricken community collectively make sense of their new, changed, and uncertain reality to adopt new survival norms [3, 27]. This includes evaluating the authenticity, source, and

extent of the changes caused by the hazard, as well as possible actions that could be taken to mitigate the risks. As the disaster continues to unfold, sustained uncertainty regarding conditions makes the sense-making process challenging; the community may suffer from information scarcity about some aspects of the disaster, while in others, they may have to sort through a vast amount of information from different sources of widely varying authenticity. Therefore, during and in the aftermath of a disaster, it is likely that members of the community collectively display an elevated level of cognitive processing. As time passes, more information naturally emerges, bringing with it increasing stability and certainty until eventually the community collectively achieves a coherent understanding of the circumstances.

Biological

Natural disasters are exogenous shocks that manifest in the physical environment of communities. Consequently, they can have direct and immediately visible effects on members of a community, including various degrees of discomfort, injury, and loss of life. Additionally, people are known to exhibit psychosomatic symptoms, such as headaches, dizziness, and soreness, due to disaster-induced psychological stress [1, 28]. The prevalence of such conditions, often medically unexplained, may indicate unresolved psychological and social issues among members of the community prompted by the disaster. Therefore, while a community may focus on immediate physical danger and harm in the short term after a disaster, there are other more long-term biological indicators that could characterize the trajectory of overall well-being.

Physical

Beyond the immediate risk of bodily harm, the most visible impact of a disaster is the destruction of infrastructure and the disruption of normal life. People in affected communities can have their homes and neighborhoods damaged or completely destroyed. Much of the attention of early disaster recovery research focused on this dimension, as it reflects the immediate goal of many large-scale relief efforts in disaster-stricken regions. Although more recent research increasingly considers the other aspects to be equal or more crucial in the recovery process, both the intensity and duration for which physical reconstruction remains in community consciousness still have the potential to influence the overall recovery trajectory of the community [3, 29]. This is because many of the elements of the physical environment (Ex: homes, utilities, and roads) correspond to the bottom tier in Maslow's hierarchy of needs [30] and as such may be pre-requisites for addressing the more complex psychological and social needs faced by members of the community.

In the rest of the paper, this thematic framework serves as the overarching organizing principle. First, it is the lens through which we identify and catalog different dimensions of community response and the corresponding linguistic markers in tweets. Second, it is also the means by which we organize and interpret our observations of how community response changes over time.

Trajectories of community response and recovery

In this section, we discuss theoretical models that shape the current understanding of how communities respond to disasters. These conceptual models often play an important role in the development of institutional practices for supporting disaster-stricken communities [6, 31]. Here, they serve as a useful baseline for placing our own observations in the context of existing knowledge on community psychosocial dynamics during disaster recovery.

In general, conceptual models of disaster response and recovery characterize these collective processes as disaster-stricken communities moving through a series of steps or arcs for simplicity while acknowledging that some of these arcs often overlap in practice [7–9, 32–34]. In these models, the community response begins immediately before the onset of a disaster, when community members become aware of the imminent danger and take protective action [32–34]. This pre-disaster stage is usually known as *warning phase*, as people actively disseminate warnings about the event within their communities, in addition to taking other protective actions. The warning phase follows a pattern of behavior that has been consistently observed both offline and online through social media [6, 35]. The degree of forewarning received by a community is likely to affect its response. Sudden and unexpected incidents are likely to evoke feelings of helplessness, fear, and vulnerability, while incidents with substantial prior warning can cause feelings of guilt and blame within a community that is not adequately prepared [3, 6]. The duration of this phase during a particular disaster depends on the level of warning the affected community received and whether the hazard, continues to evolve, while it is active.

The *warning phase* is followed by *impact* or *shock*, the period during which the hazard remains active within the affected community. Depending on the severity of the hazard, communities can experience a variety of psychological effects ranging from anxiety and vigilance to collective panic. During this period, people focus on their personal safety and that of their family [1, 3, 36]. The impact phase can last for as long as a few minutes, as in the case of an earthquake, or hours or days, as in the case of storms and wildfires.

In theoretical models, the period following the onset of disaster is variously known as the *heroic phase*, the *honeymoon period*, the *therapeutic community*, etc. [8, 32, 33]. During this stage, people are driven by a stronger prosocial orientation and engage in relief efforts within their neighborhoods. In turn, their efforts evoke a strong sense of community spirit that may last weeks. During the first weeks after a disaster, this dynamic may suppress the negative psychological and psychosocial effects of the disaster.

Eventually, people exhaust their optimism and their capacity to support other community members. Concurrently, both external attention from the wider public and relief efforts begin to wane. This support vacuum marks the end of the *honeymoon* stage. At this point, affected communities take stock of the long-term repercussions of the disaster and the challenges to returning to normal. This period is known as the *disillusionment phase* [6, 33] or *Inner Contradiction and Crisis phases* [34] in reference to the cognitive dissonance that communities experience as they

return to the harsh post-disaster reality after being buoyed temporarily by external support and social cohesion.

The disillusionment phase can last for many months [6]. Communities experience a resurgence in negative affect, as well as maladaptive behaviors, such as psychosomatic symptoms and substance abuse. Eventually, communities enter the *reconstruction* or *recovery* phase, where they come to terms with their new reality and return to stable behavioral patterns, which may or may not be similar to pre-disaster patterns.

Social media-based disaster research

Social media is widely used during and after disasters by affected individuals, media, relief agencies, and volunteers to search for and disseminate information and to coordinate relief work [37–39]. Previous Twitter-based research has covered a variety of topics, such as the extraction of disaster-relevant content, situational awareness, and information seeking and diffusion behaviors of affected communities and their collective emotions [40].

Most of these previous studies have considered the information produced and shared on Twitter during disasters from the perspective of how it serves to improve awareness of the evolving circumstances among affected communities, the general public, and relief organizations. These analyses are interested in the types of information shared and their temporal variation. Information categories such as warnings and advice, donations, and volunteerism feature prominently during disasters and their aftermath [16, 40, 41]. Although such categories provide useful contextual knowledge about Twitter use during disasters and how it can be used to improve relief efforts, they are less useful for understanding community well-being over time.

Less frequently, studies have explored the behavior of Twitter users affected by a disaster through comparative pre-disaster vs. post-disaster analysis [42, 43] as well as the evolution of behavior over time [16, 44]. However, work that differentiates behavior during qualitatively different periods (e.g., before, during, immediate aftermath, and long-term) of the disaster's life cycle is rare. Work that draws a clear contrast between the community that is affected by an event and the geographically diffused population of other interested individuals is similarly scarce. Kogan et al. [11] showed a useful case study of retweet behavior during Hurricane Sandy that incorporates both of these elements. They compare the behavior of users who were present in the affected region with that of others who simply tweet about the disaster. Further, they consider the variation in behavior in discrete time intervals before, during, and after the disaster. By contrast, we focus only on the behavior of vulnerable users over time as a continuous, evolving process.

Some studies have used social media to explore the psychological effects of disasters on affected populations. This literature has established that linguistic markers in tweets can reveal the psychological states, interests, and opinions of individual users and communities exposed to disasters [15, 16, 45]. However, these studies either focus on a single emotion, such as fear or grief, or consider the general prevalence of different emotions during the entire observation period. In contrast, we analyze

the temporal evolution of the collective psychological state of communities along multiple affective and social dimensions. One of the first studies in this space, De Choudhury et al. [45] examine the affective response and the resulting desensitization of communities in the context of the “Drug War”-related violence in Mexico. In a similar vein, Wen and Lin [46] estimate the prevalence of anxiety, anger, and sadness over time in tweets produced by a community affected by a terror attack. Furthermore, Garcia and Rimeé [47] analyze collective emotions and language related to prosocial behavior and solidarity in French tweets after the Paris terrorist attacks of November 2015. However, their work is limited to the analysis of a single event and does not identify patterns of community response that generalize across a range of disasters. In this sense, the study by Saha and De Choudhury [48], which used Reddit to investigate the temporal dynamics of stress and other psychological markers in 12 cases of college gun violence, is more closely related to our work, although it is limited to a few events of a different type of crisis.

Disaster dataset

In this section, we describe the selected set of disasters, the procedure to define spatio-temporal boundaries for disasters, and Twitter data collection.

Spatio-temporal boundaries

We use disaster declarations reported by the Federal Emergency Management Agency (FEMA)¹ between July 2011 and September 2016 to estimate spatio-temporal boundaries for reported disasters. There are 530 declarations during this period. We extract six properties from each declaration: (i) *Disaster Type* (e.g., *Flood*), *Start Date*, *End Date*, *State*, *County*, and *Aid Amount*. Since FEMA makes regional declarations when the required criteria are met, we find that multiple declarations are frequently associated with a single event. Since we are interested in studying communities as a whole, we group FEMA declarations that correspond to the same *disaster event*. Our intuition for grouping a set of declarations into a disaster event is that they must: (i) be of the same incident type, (ii) be spatially close, and (iii) be temporally close. Our grouping procedure consists of two steps:

1. **Group by type and time** Declarations are split into ordered sequences of the same type, so that the starting dates of consecutive declarations in a sequence are no more than 5 days apart. We chose the threshold of 5 days through trial and error, starting with 1 day and increasing 1 day at a time. For each threshold, we manually inspected a sample of the generated candidate groups and estimated how often and accurately their spatio-temporal boundaries aligned with actual events.

¹ <https://www.fema.gov/openfema-data-page/disaster-declarations-summaries-v2>.

- 2. Identify spatially connected sequences** For each sequence of declarations, we construct a network from the corresponding states, where each state is a node, and there is a link between two states if they share a border. Each connected component in this network is assigned to a different disaster event. We assume that disaster events do not move directly between states that do not share a border. Note that this definition allows disaster events to span large geographic areas as long as they move across connected states.

For each identified disaster, we estimate the temporal and spatial boundaries as follows.

- *Temporal boundaries* The period from the start date of the first declaration to the end date of the last declaration.
- *Spatial boundaries* First, we identify the smallest rectangle that covers all counties recorded for a particular disaster. Next, we expand the boundary in all directions by 10% of the length of the rectangle to include potentially affected peripheral regions not covered by the declarations.

Validation

This procedure produces a set of 350 candidate disasters. To evaluate their validity, we verify the dates and location of disasters using Google Trends² and online news articles. Google Trends provides a time series of the relative Google search volume for a query during a given time interval and geographic area. For each group, we search for the disaster type, restricting to the temporal bounds of the disaster and the states involved. We verify that the search volume peaks around the time of the first declaration for most candidates. In cases where the search volume is unavailable, we search for news articles about the disaster to confirm the validity of our groups. When our best estimate for the start of a disaster based on this approach differs from the first declaration of the corresponding group, we use our estimate. The starting dates of 37% of the disasters are adjusted this way with mean and median adjustments of 5.6 and 3 days, respectively.

Social media data

We use the following procedure to identify Twitter users associated with each disaster and collect their tweets.

Disaster-affected Twitter users

We use an archive of 10% of the Twitter stream to identify users who were in the affected area during the time of the disaster. We define affected users as those who


² <https://trends.google.com/trends/>.

had at least one geo-tagged tweet within the spatial boundaries of each disaster between 1 week before the start of the disaster and 1 week after the end of the disaster. Of the 350 disasters identified from the FEMA declarations, we select 203 disasters that had at least 100 affected users. This dataset contains 13 different disaster types; Severe Storm(s) (72), Fire (66), Flood (35), Severe Ice Storm (7), Hurricane (5), Snow (5), Tornado (3), Other (3), Coastal Storm (2), Earthquake (2), Mud / Landslide (1), Typhoon (1), and Volcano (1). The total number of users identified this way is 3,984,530.

Collecting Twitter user timelines

For each disaster, we focus on a period spanning 8 weeks, from 4 weeks before the start date to 4 weeks after it. We collect the 3200 latest tweets (a limit imposed by Twitter), including retweets, for each user in our sample using the Twitter API. Approximately 8% of the users in all disasters had more tweets than this number at the time of collection. For some of these users, particularly very active ones, we are unable to collect all tweets during the disaster period. In addition, some users have made their tweets private or deleted their accounts. Finally, since our goal is to capture changes in the language of users before and after the disaster, we limit the analysis to users who posted at least once before and once after the onset of the disaster. This ensures that our baseline observations before the disaster are generated by the same population of users as those we observe after the disaster. We are able to obtain a total of 2,082,210 users who meet these criteria corresponding to approximately 227 million tweets. Figure 1 shows the distribution of the number of users across disasters after applying this final constraint. Figure 2 shows that disasters in our data set are spread over a wide range of geographic areas in the United States.

Fig. 1 Histogram of disaster size in terms of the number of associated Twitter users






Fig. 2 Data volume maps in the contiguous U.S. Left: number of incidents per state in our dataset. Right: aggregated number of users appearing in all disasters in each state

Methods

We begin this section with a discussion of the rationale for choosing a 4-week period after the onset of disaster for our study. Next, we describe our methodology to characterize community responses. We use a lexicon-based approach to estimate different dimensions of community behavior based on tweets from affected users. We use these estimates to build a multidimensional time series of disaster response for each disaster that controls for pre-disaster behavior. Then, we cluster these responses to identify groups of disasters that have similar trajectories along different dimensions. Finally, to discover any prototypical responses, we compare the contribution of dimensions to the differences between clusters.

Length of the study period

Two important methodological questions related to the study period must be answered when quantifying and comparing the community response in many disasters. The first is whether we should vary the duration of the study period for each disaster, since they may differ in terms of the duration of immediate impact. In this study, we choose to fix the length of the study period for all disasters, since having different study periods across disasters would make any form of temporal comparison challenging. We conduct a post-comparison statistical test to establish if our findings are affected by a relationship between the duration of a disaster and the community response.

The second question is determining a study period that is appropriate for all the events. In this study, we choose a 4-week period after the onset of the disaster for two reasons. First, due to normal attrition of Twitter users, the dataset has fewer users, and thus, confidence in observed community behavior would be lower further away from the onset of disaster. Second, a longer study period increases the probability that a community will be exposed to new disasters or other shocks that would interact with their ongoing response to the original disaster.

Linguistic dimensions of disaster response

We estimate the aggregate prevalence of linguistic markers associated with different themes of response for each disaster over a 2-month period. A supervised learning approach to achieve this goal would have needed a large volume of human-annotated tweets. Instead, we use a lexicon-based approach. We use lists of words associated with the markers selected from validated existing lexicons to estimate a normalized index for each linguistic marker by counting the number of associated words in tweets. Previous work has shown that when using established lexicons, this approach is effective in estimating psychological and social markers in text, including tweets [49, 50]. Further, recent studies comparing the performance of lexicons against large language models (LLM) have shown that linguistic time series constructed from tweets using both methods correlate well with data from representative surveys and that, in this regard, lexicon performance is comparable to LLM [51, 52].

We use LIWC (2015)³ as the primary lexicon. One limitation of LIWC is that it contains relatively few categories that focus on emotion (*positive, negative sentiment, anger, anxiety, and sadness*). One of the often mentioned affective dimensions in the literature on disaster response is *fear*. Since it is absent in LIWC, we use the corresponding word list from the NRC lexicon [53] to supplement our analyses with the dimension *fear*. Table 1 shows the list of dimensions used to represent the response to disasters (22 from LIWC, 1 from NRC).

Disaster response trajectories

For each disaster, we estimate an aggregate time series for each dimension. Let N_i be the number of tweets on day i of the disaster and $N_{p,i}$ be the number of tweets on day i that contain at least one word of dimension p . The measurement of category p on day i of the disaster is $T_{p,i} = \frac{N_{p,i}}{N_i}$. Next, we adjust the estimates to account for any differences in the content of people's tweets due to what day of the week it is. We achieve this by fitting an additive time series model to the data and subtracting the estimated day-of-the-week seasonal estimates from our observations.

These numbers include retweets, replies, and quote tweets in addition to original tweets. We include retweets, since the act of retweeting suggests that any affect in the tweet is shared between the original user and the retweeting user. This is particularly relevant in our case, as previous work has shown that when a disaster strikes, affected users preferentially retweet others from the same community [11]. In the case of replies and quote tweets, we include only the text of the replying/quoting user.

Each disaster in our dataset affected a different community defined by unique geographical boundaries. The typical behaviors of these communities can differ for a variety of social, cultural, demographic, and economic reasons. To compare the response of the community across disasters, we need to account for these differences

³ <https://liwc.wpengine.com/>.

Table 1 Selected lexicon dimensions

Theme	Lexicon dimension	References
Psychological	Anxiety	[6, 54]
	Anger	[6, 17]
	Swear	[6]
	Sadness	[6, 54]
	Fear	[6, 54, 55]
	Achieve	[6, 54, 55]
Psychosocial	Reward	[6, 55]
	Social	[6, 29]
	We	[6, 56, 57]
	Family	[3, 6]
	Friend	[3]
Sense-making	Affiliation	[6, 57]
	Risk	[6, 29, 55]
	Cause	[55, 56, 58]
	Tentative	[55, 58]
	Certain	[55, 58]
Biological	Differ	[55, 58]
	Insight	[56, 58]
	Health	[6, 28, 54, 56]
	Ingest	[28, 55, 59]
Physical	Body	[6, 28, 54]
	Home	[3, 6, 29]
	Work	[3, 29]
	Money	[3, 6]

All dimensions except Fear, which was created using the NRC lexicon, are sourced from LIWC

in the typical pre-disaster behavior of communities. We do this by adjusting observed behavior in the temporal vicinity of the disasters by typical behavior uniquely estimated for each community from a corresponding pre-disaster baseline period. For each dimension p , we split the time into a baseline period of 3 weeks, beginning 4 weeks before the start of the disaster and ending 1 week before it, and a study period of 5 weeks from 1 week before the start to 4 weeks after it. Next, we construct a time series for the study period defined by the relative difference in intensity $T_{p,i}$ for each dimension p during the study and the baseline periods. More precisely, we estimate the baseline rate for dimension p , P_p , as the fraction of tweets that had at least one word from p from all tweets during the baseline period. The normalized estimate for p on day i^{th} of the study period is defined as $N_{p,i} = \frac{T_{p,i} - \mu}{\sigma}$, where μ and σ are the mean and standard deviation of the distribution of daily estimates during the baseline period.

Table 2 Distribution of events by the type of hazard in Cluster 1 and Cluster 2

Type	Cluster 1	Cluster 2
Severe Storm	24	48
Fire	24	42
Flood	5	30
Severe Ice Storm	1	6
Snow	0	5
Hurricane	0	5
Other	1	2
Tornado	0	3
Earthquake	2	0
Coastal Storm	0	2
Mud/Landslide	0	1
Volcano	1	0
Typhoon	0	1

Clustering response trajectories

Our first objective with this study is to establish the existence (or absence) of regular patterns in the community response in different disasters. We approach this objective using a clustering approach. Clustering can highlight the aspects of community response over social media that differ across events substantially as these dimensions will contribute to the decision criteria that define the resulting clusters. Conversely, the dimensions that do not contribute to the clustering decision are those that have similar trajectories across the clusters (i.e., all disasters). Next, we describe how this rationale is extended to identify prototypical patterns of community response, as well as elements of response that show marked differences between groups of disasters.

The adjusted community response for each disaster is represented by a 5-week, 24-dimensional time series. Our next goal is to find clusters of disasters with similar response time series. The estimated time series suffers from noise to which most well-known time series clustering methods are sensitive, particularly under high dimensionality [60]. Therefore, we first smooth individual dimension time series by fitting locally weighted first-order regressions with a span of 20%. Next, we estimate the pairwise dissimilarity between response trajectories using the multidimensional dynamic time warping distance (DTW) [60]. DTW is generally superior to Euclidean distance in the case of time series clustering, since it is purpose-built to account for temporal shifts: it finds the best compromise between alignment of the shapes of two time series and the cost of shifting (in time) necessary to achieve a given alignment. We use hierarchical clustering with Ward's criterion on the resulting trajectory–trajectory distance matrix to separate our dataset into groups of disasters that have similar community response trajectories. We evaluate the quality of the results for different numbers of clusters ranging from two to ten using Silhouette Coefficient, Dunn Index, and Connectivity (C). We find that dividing the disasters into two

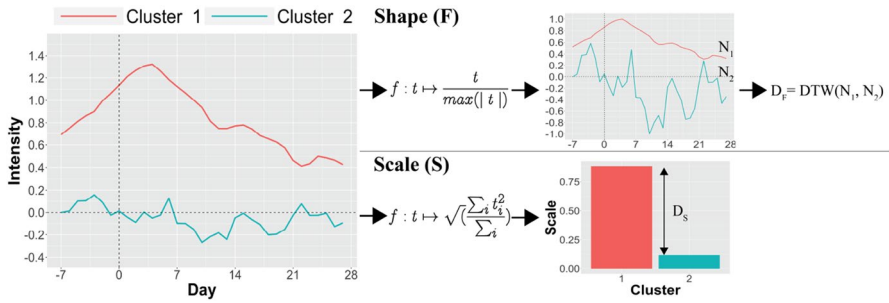


Fig. 3 Scale and shape difference for *anger* between Clusters 1 and 2. The difference in scale is measured as the absolute difference in the root mean sum of squares of magnitudes from the average trajectories of the two clusters. The shape contribution is measured as the DTW distance between magnitude-normalized average trajectories

clusters provides the best results based on all three metrics (Fig. 8). Disasters are unevenly divided among the two clusters, the smaller cluster (Cluster 1) containing 58 disasters and the larger one (Cluster 2) containing 145 disasters. Table 2 shows how different types of disasters are distributed between the two clusters. We generate a representative signature of each cluster using DTW barycenter average (DBA), a method that finds a representation that minimizes the sum of the squares of the DTW distance to the members of a cluster [61].

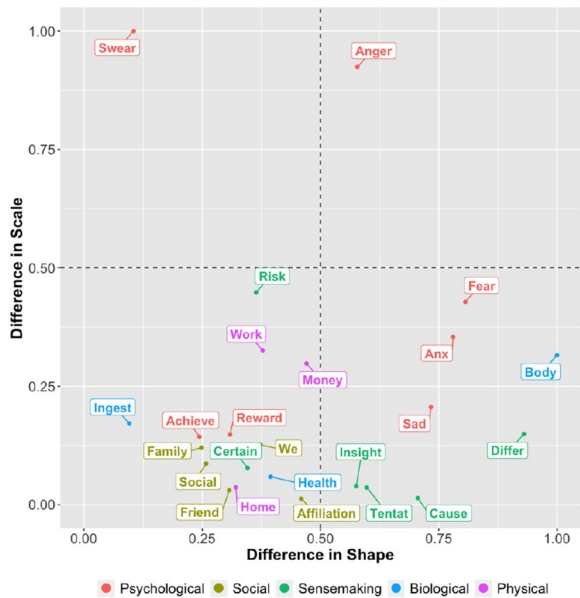
Results

Next, we discuss our findings in terms of disaster response trajectories and address how they relate to our research questions. Note that we base our observations and consequent interpretations solely on the trajectories relative to baseline period. This is true even for cases where we compare the trajectories of different linguistic dimensions with each other. As an example, if, on a particular day, the prevalence of *Fear* in a community is one standard deviation above its mean during the baseline, while that of *Social* is only half a standard deviation above its corresponding mean, we would interpret this observation as *Fear* being more prominent within the community than *Social*, even if, in absolute terms, *Social* is much more prevalent in community tweets.

RQ1: universal patterns in disaster recovery trajectories

While the two clusters correspond to distinct disaster response patterns, we find that the differences between them are not evenly distributed across the different dimensions. Further, the differences between clusters along a given dimension can stem from two sources; (i) scale of observations and (ii) shape of the time series. The scale of observations signals the intensity of the community experience, whereas the shape of the trajectory is more informative of how behavior evolved over time. To disentangle these two aspects of scale and shape, we

Fig. 4 Scale vs. shape contributions to differences in dimensions. The values in both axes are scaled to the range of [0,1]



estimate them separately, as shown in Fig. 3. For a given cluster, we estimate the characteristic scale of a dimension as the root-mean-square magnitude of its average trajectory. Consequently, the difference between the two clusters is the difference in the characteristic scales. To estimate the difference between clusters in terms of just shape for a given dimension, we normalize the average trajectory of each cluster by the largest magnitude observed in it. This normalization scales a time series to approximately the same scale with a maximum magnitude of one, while preserving its shape. Then, we estimate the difference in shape between the two clusters as the DTW distance between these normalized average time series. Figure 4 shows how the differences in dimensions between clusters are distributed between shape and scale. For dimensions above the 45° line, the difference across clusters is driven more by the scale of observations than the shape, while the reverse is true for those below it. We consider dimensions that fall within the lower left quadrant of the figure, defined by (Scale < 0.5 and Shape < 0.5), to be dimensions for which the clusters have similar trajectories. Thus, 14 dimensions with prototypical patterns: *achieve*, *reward*, *social*, *we*, *family*, *friend*, *affiliation*, *risk*, *certain*, *health*, *ingest*, *home*, *work*, and *money*. Figure 5 shows the average time series in all disasters for these dimensions. Next, we describe our findings based on these average time series and discuss how they relate to the existing theory on disaster response.

Sense of productivity is lost

The normal functioning of a community depends on the stability of its physical and social environments. Disasters tend to suddenly and significantly disrupt these. In

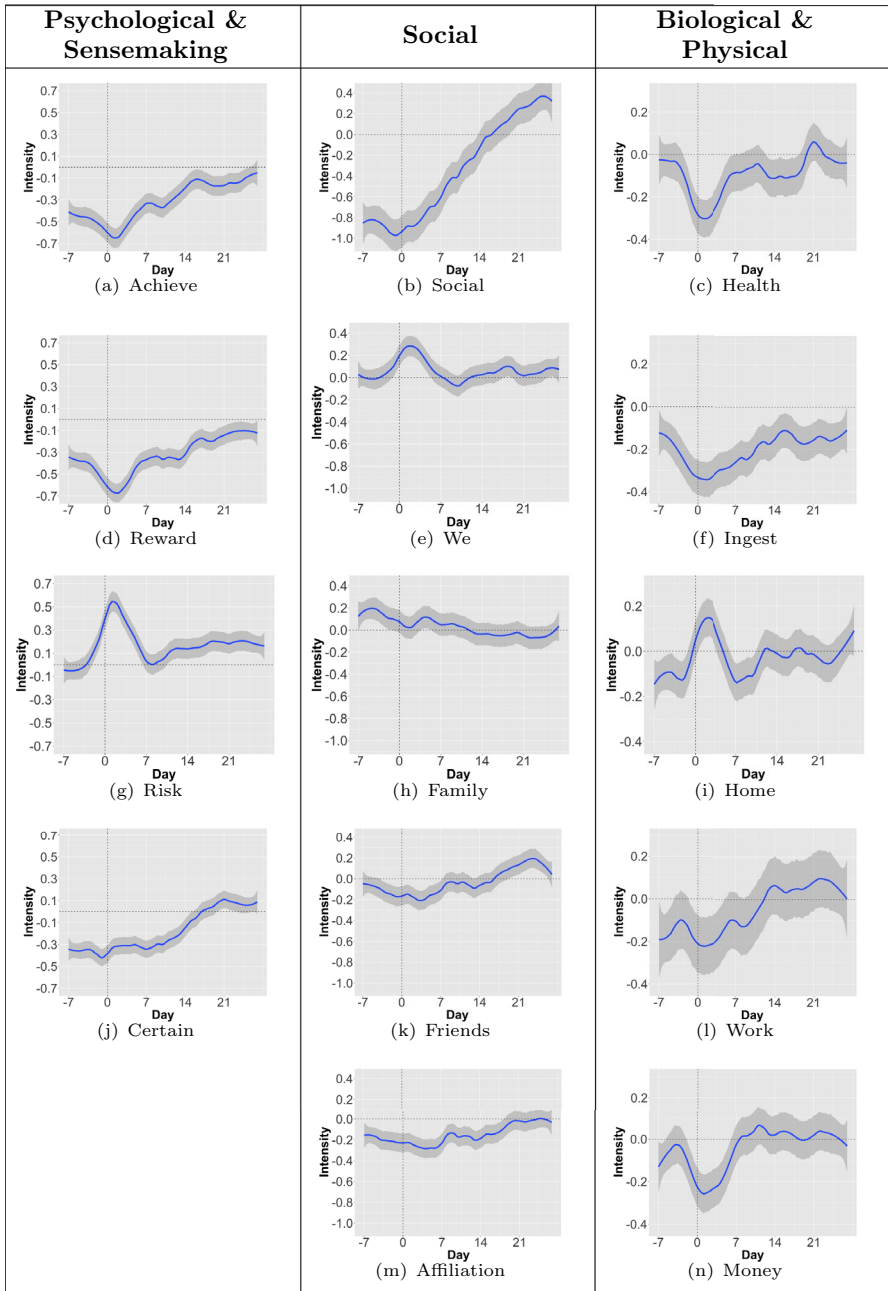


Fig. 5 Dimensions with universal trajectories across disasters in the dataset. All figures in a column have the same intensity (Y) scale, and the shaded regions are 95% confidence intervals

our results, the two prototypical dimensions of sense-making, *Certain* (Fig. 5j) and *Risk* 5g, reflect how communities experience this disruption. We observe that communities experience a high level of uncertainty from before the onset of the disaster up to 2 weeks after it, before gradually recovering to normal levels. In comparison, communities only perceive a significant threat to themselves immediately after the event while there is an active hazard which is reflected as a spike in *Risk*. A more detailed analysis shows that the increase in *Risk* is mainly driven by tweets associated with warnings (Fig. 9).

In addition, we observe that the two psychological dimensions, *Achieve* (Fig. 5a) and *Reward* (Fig. 5d), have nearly identical prototypical patterns. They show a significant decline from baseline in the immediate aftermath of the disaster followed by a gradual return to normal by week 4. To unpack this observation, we examine the contribution of individual words to these categories (Figs. 10 and 11). In *Achieve*, the decline is driven by words that imply celebration (e.g., win, celebrat*, proud, award) or a positive outlook in general (e.g., improv*, challenge*, achieve, earned). The decline in these words is counterbalanced to some extent by an increase in comparatively neutral words (e.g., work, team*, lead). *Reward* follows a similar, if less clear, pattern with words that imply a positive outlook (e.g., good, goal, scor*, promot*, adventur*) declining, while more prosaic words (e.g., get, got, take, earn*) become more frequent. There are several possible behavioral changes within affected communities that, individually or in combination, explain these observations. In the aftermath of the disaster, when the devastation is easily apparent, people may feel less celebratory. Another possibility is that the disruption caused by the disaster reduces people's engagement with the spheres of life they associate with goals and a sense of achievement (e.g., Career, School, Sports). A third possibility is that people in the affected community are being more selective with their Twitter use, choosing to focus on disaster-relevant content and self-censoring any routine positive content that they deem irrelevant or inappropriate. We leave the disentanglement among these possibilities for future work.

Social life contracts and gradually relaxes to pre-disaster levels

Given the disruption to local social life that can be caused by disasters, it is possible that communities would show a reduced level of social activity—at least for a while—after a disaster. However, previous work suggests that people would demonstrate greater community spirit and pro-sociality as they mobilize collectively to face the effects of a disaster.

Among the prototypical patterns in our results, *Social* (Fig. 5b) dimension, a measure of overall sociality, shows a substantial decline from baseline at the start of the disaster and gradually recovers to normal levels by week 4. Although this supports the first argument, a look at the contributions of individual words to this aggregate observation provides valuable nuance. In the aftermath of disasters, the observed decrease in *Social* is mainly due to a reduction in the relative rate of

retweets (Fig. 12).⁴ Communities tweet and retweet more often during and after a disaster, but are less likely to retweet compared to normal times. We suggest that the relative decline in retweets is due to affected communities becoming more discerning with the content they boost with their limited energy. Previous work has shown that vulnerable populations focus on locally produced content rather than the much larger volume of outside content during a disaster [11].

Food-related processes rebound slower than other aspects of life

We observe that among the dimensions associated with biological and physical processes, disaster-induced perturbations have subsided quickly in all except one dimension, the *Ingest* dimension associated with food-related behavior. For dimensions *Health*, *Money*, and *Work*, the onset of disaster results in a decline in attention, but levels return to pre-disaster levels within a week. A word-level analysis of this short-lived decline shows that it corresponds to affected communities referring to routine preexisting behavior and conditions less often. For example, under *Health*, people were less likely to mention the words doctor*, fitness*, cancer*, and diet* among others (Fig. 13) and under *Money*, words, such as market*, shopp*, sale, and debt*, occurred less often in their tweets (Fig. 14).⁵

Compared to other dimensions that capture biological/physical processes, *Ingest* takes much longer (4 weeks) to return to pre-disaster levels from the perturbation (decline) observed around the onset of disaster. A word-level analysis shows that the decrease in *Ingest* is driven by words, such as dinner, lunch, beer, cake, and coffee (Fig. 17). Other words, such as water, eat, hungry, egg, rice, and thirsty, actually increase in the aftermath of disaster.⁶ Words that decrease in frequency appear to be associated with relatively high-level food needs and activities, while those that increase correspond to references to basic sustenance [62]. We suggest that this observation reflects the reality that disasters inflict substantial damage to food infrastructure (restaurants, supermarkets, and grocery stores), limiting food diversity and accessibility within affected communities. It is more difficult to address why food-related activities take longer to rebound compared to those related to health, employment, and finances. One possibility is that food systems rely on relatively more types of other critical infrastructure (such as energy for equipment, water for growing, roads for transportation, and communication technology for coordination) that are likely to be disrupted by disasters [63]. In fact, due to this dependence, the state of the food infrastructure has been suggested as a possible indicator of short- and long-term recovery from disasters.

⁴ The text of a retweet contains the prefix "RT" when it is recovered from the API.

⁵ These observations were consistent between the entire post-onset period as well as considering only the first week after onset.

⁶ We also observe a substantial increase in smoke*. However, we associate this increase with wildfire smoke as opposed to the culinary practice of smoking.

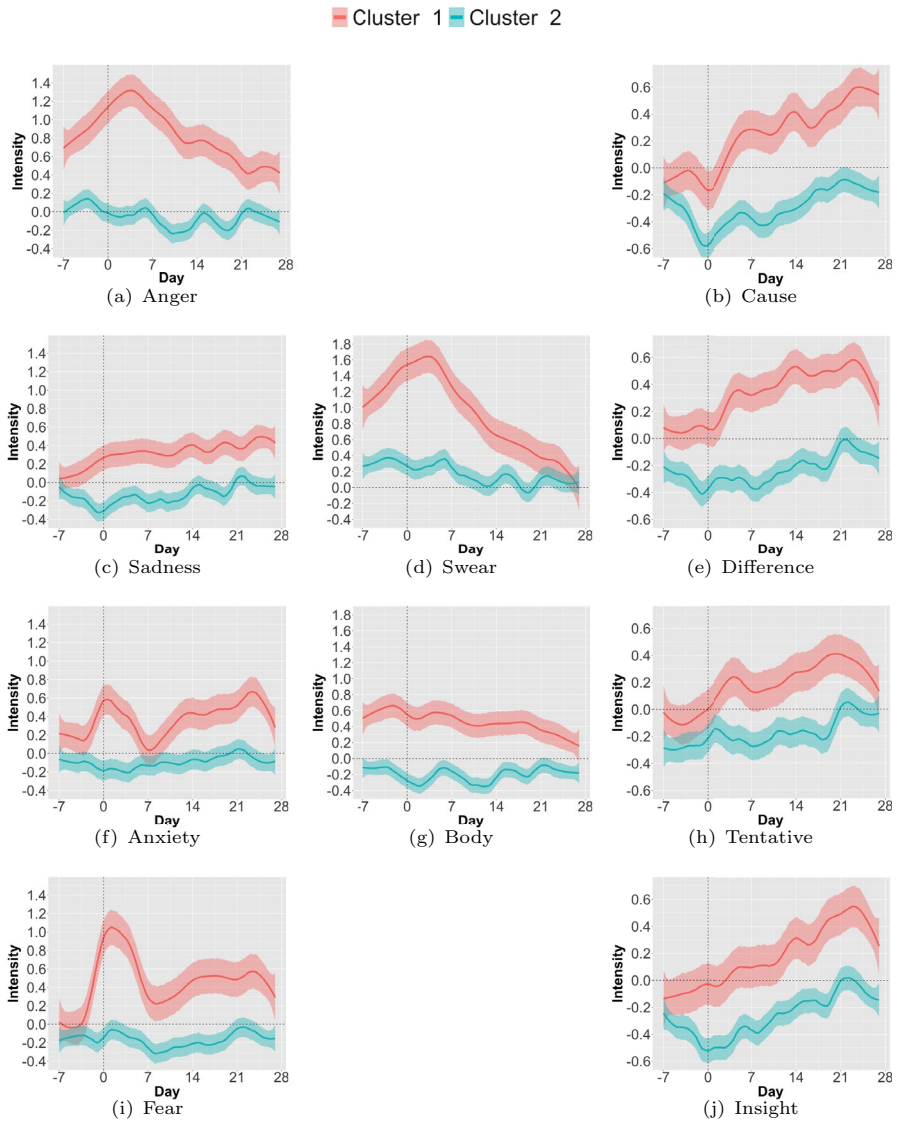


Fig. 6 Categories with noticeable differences across the two clusters. All figures in a column have the same intensity (Y) scale. The shaded regions are 95% confidence intervals

RQ2: variations in community disaster response

We identify dimensions that have large differences between clusters as those that fall outside the lower left quadrant (Scale > 0.5 or Shape > 0.5), which includes all affective psychological dimensions (*anger*, *swear*, *sadness*, *anxiety*, and *fear*),

Table 3 Odds ratios by covariate from logistic regressions predicting the cluster associated with a disaster ($*p < 0.05$)

Model	Scale(p)	Duration(p)	Storm(p)	Fire(p)	Flood(p)
Type	–	–	1.07 (1)	1.22 (1)	0.36 (1)
Duration	–	0.88 (0.465)	–	–	–
Scale	1.17 (0.313)	–	–	–	–
Type + Duration	–	0.87 (0.447)	1.05 (1)	1.25 (1)	0.36 (1)
Duration + Scale	1.22 (0.223)	0.83 (0.329)	–	–	–
Type + Duration + Scale	1.44* (0.046)	0.79 (0.218)	1.03 (1)	1.48 (1)	0.29 (1)

four sense-making dimensions (*insight, cause, differ, and tentative*) and the physical dimension of *body* (Fig. 6).

Overall, the two disaster clusters are differentiated by a consistent difference in the magnitudes of the time series. The smaller cluster of 58 disasters (**Cluster 1**) has markedly higher intensity at one or more points during the period across all dimensions that differ between clusters. In comparison, the larger cluster of 145 disasters (**Cluster 2**), with the exception of *swear*, remains at pre-disaster levels or drops below it. In relation to prior work, we hypothesize that the smaller cluster may correspond to more severely affected communities leading to a strong emotional response and a need to engage in sustained sense-making [6, 64]. In the following sections, we first investigate if the differences between the two clusters may be explained by intrinsic properties of the disasters. Then, we describe the main themes that emerge from the temporal patterns in Cluster 1.

Effect of disaster characteristics on community response

In discussing the findings, we refer to potential relationships between response trajectories and disaster characteristics, such as the type, scale, duration, speed of onset, level of forewarning, and familiarity of the disaster drawn from existing disaster research [3, 6, 65, 66]. Next, we evaluate the association between three of these properties—type, scale, and duration—and the observed clusters of disaster response.

We derive the values for each property of a disaster in the following way. First, in terms of the type of disaster, severe storms, fires, and floods account for 85% of the 203 disasters used in this study. Therefore, to mitigate the effects of sampling bias, we limit the investigation to only disasters of these three types. Cluster 1 contains 24 storms (42.25%), 24 fires (42.25%), and 5 floods (9.5%), while Cluster 2 includes 48 storms (40%), 42 fires (35%), and 30 floods (25%). Second, we consider the total amount of federal assistance received by communities through FEMA as a proxy for the scale of each disaster. Finally, we consider the number of days between the start date and the end date of a disaster according to the FEMA declaration as its duration.

We use a series of logistic regressions that consider all different combinations of properties to investigate the association of these properties with the membership of the disaster in Cluster 1 or 2. For this purpose, we represent the type as three separate indicator variables (Storm, Fire, and Flood). As shown in Table 3, the only statistically significant relationship is between the scale of the disaster and the cluster in the complete model that controls both the duration and the type. This suggests that once the type and duration are taken into account, an increase of one standard deviation on the scale increases the odds of a disaster belonging to Cluster 1 by 44%. This provides some support for our hypothesis that Cluster 1 corresponds to disasters that have a greater impact based on the observation of trajectories.

Theme 1: fear and uncertainty recur after a delay (Cluster 1)

In Cluster 1, we observe a zig-zag pattern of increase and return to baseline levels in *fear* (Fig. 6i), *anxiety* (Fig. 6f), and *tentative* (Fig. 6h). These dimensions initially peak at the onset of disaster or in the immediate aftermath, but quickly return to pre-disaster levels within the first few days, only to increase again almost immediately. Communities maintain an elevated levels of fear, anxiety, and uncertainty in remaining 2 weeks within the study period. Interestingly, this subsequent peak in the trajectory is higher than the original for *anxiety* and *tentative*, but lower for *fear*. What may be driving this behavior within the communities of Cluster 1?

Shortly after the onset of a disaster, members of communities come together in solidarity and experience greater social cohesion, a period sometimes known as the *honeymoon* phase in prior research [6, 33]. We suggest that this process temporarily mitigates some of the negative emotional effects of the disaster. The end of the *honeymoon* phase is marked by a decline in attention and support outside of the community, as well as the gradual exhaustion of social support within. During this period, communities gradually shift their attention to the long term and realize that they have many challenges ahead on the road to recovery, which can lead to a resurgence of concern. The fact that expressions of *fear* within these communities do not reach the same levels of intensity as observed at the onset of disaster, whereas *anxiety* and *tentative* exceed them, may be simply due to people being more likely to respond to the immediate dangers in the aftermath of disaster with fear. In contrast, we suggest that uncertainty associated with the long term is more likely to elicit expressions of worry or concern.

Theme 2: communities increasingly engage in sense-making in the aftermath (Cluster 1)

Cluster 1 communities show a pattern of gradual increase in intensity in the cognitive processing dimensions *cause* (Fig. 6b), *difference* (Fig. 6e), *tentative* (Fig. 6h), and *insight* (Fig. 6j), starting from the immediate aftermath of disaster to nearly the end of the study period. In *difference*, *tentative*, and *insight*, this pattern reverses during the last few days of the study period, which may be an early sign that these

communities are settling on a *new normal*. The fact that *insight*, which is associated with expressions of understanding, lags behind the other three dimensions and only peaks in the final week of the study period support this argument.

We suggest that the presence of elevated levels of cognitive processing in Cluster 1 communities, as well as the absence thereof in Cluster 2 communities, indicate that these communities experience greater upheaval and have more difficulty facing the challenges brought on by disasters. This may be at least part due to Cluster 1 being associated with more destructive disasters. However, it also suggests that these communities are more likely to have lacked the institutional, financial, and social resources to weather the aftermath.

Limitations

We acknowledge some limitations in our work. First, the Twitter data for each disaster may not be representative in terms of demographic or socioeconomic groups within the communities. We begin with a 10% sample of all tweets and then use geo-tagged tweets to identify affected users. Geo-tagged tweets represent only about 1–2% of all tweets [67]. Users who geo-tag tweets are more likely to be young, more wealthy, reside in urban areas, and belong to certain ethnic groups [67]. While our study does not adjust for these biases, the fact that our many of our findings align with the existing theoretical models of offline community experience suggests that deviations between different demographic and socioeconomic groups are more likely to be in scale rather than in form.

Second, we have limited our analyses in this study to an observation period that extends only 4 weeks beyond the onset of the disaster. We opted for this short-term analysis to manage computational complexity as well as the difficulty of disentangling responses associated with individual disasters, when communities, or parts of them, experience multiple disasters or other relevant events over time. Prior theoretical and empirical work suggests that recovery from disasters can take years. It is uncertain if our methodology of measuring Twitter trajectories and the community posts as a source of data are adequately sensitive and precise to measure community behavior over longer time periods.

Third, pre-disaster characteristics of communities, such as social and economic capital, inequality, and demographics, influence community resilience to the adversities of disaster [1, 3, 4]. How these community properties interact with the type, magnitude, and duration of disasters to produce different recovery patterns is an interesting research question for future study.

Finally, in using aggregate community-level estimates for social and affect categories of response, we implicitly assume that a community is monolithic. In reality, members of a community may demonstrate a range of responses. There is a risk that the collective behavior observed in our study is distorted due to individual responses not being amenable to linear combination. This limitation can be resolved in future

work in two ways: (i) conducting the analysis at a more spatially granular level, such as a county or census tract level, where demographic and socioeconomic data are readily available and (ii) employing statistical techniques such as latent growth mixture modeling to automatically extract different common recovery trajectories that play out within a single community experiencing a disaster.

Discussion

In this study, we demonstrate that community response to disasters over social media, as measured by lexicon-based time series, follows regular temporal patterns that generalize well to different disasters and communities.

This study demonstrates that lexicon-based time series derived from the social media posts of people in disaster-stricken communities can effectively capture many theoretically expected phenomena in the collective process of recovery. In particular, we find that there are broad similarities in community social media behavior in terms of some aspects of their lives (e.g., social, biological, aspirational, and physical) despite significant variability in terms of the type, scale, duration, level of forewarning, etc., of the disaster. However, the results also reveal that around a quarter of the communities demonstrate trajectories that diverge from the majority along some aspects of their experience. In particular, this minority's social media conversations show sustained levels of fear, anger, and anxiety, as well as signs of an elevated cognitive load associated with sense-making; both of which are signs of ongoing challenges to adapting post-disaster *normal*. In this section, we discuss the implications of this study and its findings for two domains: (i) disaster relief and recovery efforts and (ii) future disaster research.

Implications of disaster relief and recovery assistance

Strategic disaster relief planning Relief organizations engaged in disaster recovery assistance face a difficult challenge. To effectively support communities in the aftermath of a disaster, they require two kinds of knowledge. First, at the strategic level, these organizations need to know how, in general, community response to disasters unfolds over time, to inform their policies for resource planning for assistance programs and training relief workers. The results of this study show that the behavior over social media of disaster-stricken communities follows a regular trajectory along many aspects of life and that, in many cases, this prototypical narrative aligns with prior knowledge of *real* (offline) community experience. Further, this comprehensive but easily communicated social media narrative is mapped to a timeline that is more precise than traditional models of community response. Therefore, we believe that our findings can complement the existing suite of planning tools of relief organizations.

Community response monitoring and forecasting Even though an idealized narrative of disaster response is useful as an organizing principle for support organizations, each disaster is a unique event resulting from a specific hazard affecting a particular community. Thus, the community response to a new disaster follows its own unique trajectory that can deviate from the prototypical trajectory. For example, we show that communities differ in the intensity of their emotional and cognitive response to a hazard as well as in the timing of their return to stable post-disaster norms. As a consequence, during the course of an assistance program, relief organizations are faced with the continuous need to reconcile facts on the ground with the idealized narrative to align their activities with the needs of the affected community. Our work furthers support for this endeavor in three ways. First, our methodology provides a rapidly deployable, relatively low-cost, and near real-time approach to get a holistic picture of a community as it responds to a disaster. In the event of the onset of a disaster, this method would involve collecting easily accessible public social media posts from members of the affected community and generating a response trajectory using our estimation technique. Moreover, since this approach accounts for idiosyncrasies of the community, these observations can be compared directly with the prototypical narrative. This leads to the second benefit: relief organizations can make early forecasts on the recovery process, including the characteristic response cluster, areas that may need more or less attention compared to original expectations (e.g., material support, mental health services, information services, and social programs), and more precise timing for crucial transitions (e.g., initial shock, to the honeymoon period, to a long-term perspective). In the supplementary section A.6, we demonstrate a proof of concept analysis to compare the forecasting capability of the response trajectories observed in this study. The results show that simply assigning an emerging community response to one of the clusters based on that community response from the onset of disaster until the present is able to provide a forecast of its future trajectory that outperforms time series forecasting methods around 75% of the time while using the same data.

Evaluating disaster relief interventions Another challenge that relief organizations face is the difficulty of assessing the impact of their interventions in the aftermath of disasters. Our methodology and findings may be adapted for this purpose in the following manner. First, starting from the onset of disaster, we construct the response trajectory for the affected community. We expect that even an early short-term fragment of this trajectory will contain quantitative markers for idiosyncrasies, such as the intensity of emotional response and the characteristic time to stable post-disaster behavior, stemming from the specific hazard and the specific community. Then, we are able to forecast future community behavior by adjusting the prototypical narrative to the observed eccentricities. Finally, we are able to measure the effect of assistance as the deviation of actual behavior from the expected after targeted interventions. In practice, this approach is limited to measuring the overall effect of assistance due to numerous overlapping relief programs rather than that of individual efforts that are normally active in a community after a disaster. Further, our

work highlights an opportunity for the technology community to address the needs of relief organizations and disaster-affected communities through the development of software tools that automate these functions.


Public perception of community response In addition to formal organizations dedicated to relief efforts, the public invariably plays a role in supporting communities stricken by disaster, either through community organizations or as individual Samaritans. However, unlike the organizations, which are invested in providing relief and have established information channels regarding the needs of the affected community, the public relies on and is influenced by the freely available—and at times incidentally consumed—information in their every day news sources. For many, the primary mode of such information is likely to social media [68]. One of the themes that emerged from the findings of this study is that, similar to the findings of prior work on offline behavior, community response to disasters over social media is a combination of multiple processes unfolding concurrently but not necessarily following the same shape or temporal order. Some processes, such as emotional upheaval and community bonding, operate on short time scales (days), while other dynamics such as sense-making last much longer. Similarly, processes like emotional upheaval and sense-making are marked by an increase in the presence of associated linguistic markers in community social media feeds, while others, such as social engagement and nourishment, are marked by their absence. Only content that exists on social media can be amplified by it and brought to a wider audience. Therefore, the differences between processes of community response are also likely to lead to a differences in the level of attention paid to corresponding community needs by the general public.

Implications for disaster research

Recent decades have witnessed an increase in the worldwide frequency of natural disasters, which are largely driven by climate change [69, 70]. Correspondingly, large-scale disasters that cause widespread economic and social damage and less observable effects, such as internal displacements, have become more common [71, 72]. These trends have added some urgency to our need to understand community response to disasters—in particular, around the questions of what makes a community resilient to disasters and how relief operations affect post-disaster recovery. In closing, we discuss the contributions of our study in this context as well as future research directions opened up by this work.

Measurement It is challenging to study community response and resilience to disasters in a widely applicable and reproducible manner due to the many complex social phenomena that simultaneously unfold along many interrelated dimensions. Traditional approaches, such as post-disaster surveys, interviews, or even ad-hoc analysis based on observational data, such as social media posts, are not suited for standardized measurement along these numerous dimensions for a wide

Fig. 7 Histogram of the number of lexicon dimensions in terms of the percentage of disasters for which we cannot reject the hypothesis of normality ($\alpha = 0.05$)



range of disasters. Interview and survey data suffer from biases due being collected after the fact. They are also costly to collect. Further, whether using the traditional or social media data, prior work has very often not accounted for the idiosyncrasies of individual events—partially because many studies consider only a single event. Thus, our most immediate contributions to disaster research are (i) a comprehensive thematic framework for measuring disaster response based on the existing work and (ii) a scalable method, built upon that framework, for quantifying the trajectory of community response and that allows direct comparison between disasters.

Universality In this study, we limit our focus to disasters in the United States over a 5 year period. A logical next step would be to investigate if our findings will generalize across different cultures and times. Will community responses to disasters remain consistent with observations over time? Do we observe the same prototypical narrative of disaster response in communities outside the U.S., in regions, such as Asia and Africa?

Resilient communities and disaster assistance Resilience is an increasingly important topic in both disaster research and policy, as high impact disasters that impose large economic and social costs become more common. A resilient community re-emerges from the crises of disaster in a manner epitomized by the German philosopher Friedrich Nietzsche's adage—"That which does not kill us makes us stronger" [73]. In the aftermath of disaster, resilient communities establish a *new normal* that is both better than the pre-disaster baseline (in terms of greater life satisfaction, social cohesion, etc.) and are consequently better adapted to face future crises. Prior work is rich in proposed relationships between community properties, such as social capital and wealth inequality, and disaster resilience. Since our methodology captures the effects of disasters across different communities in a consistent manner, it can serve as the foundation for principled future investigations that relate properties of communities to their resilience to disasters. Further, we expect that it will have a similar utility for analyses that evaluate the effectiveness different assistance activities in improving disaster recovery outcomes of communities.

Appendix A


A.1: Normality of pre-disaster variation in lexicon dimensions

We standardize the intensity time series for each linguistic dimension during the 35-day (5-week) study period with the mean and standard deviation in intensity estimated from the 21-day (3-week) period immediately prior to the study period. This normalization is meant to make time series for each dimension comparable across all disasters. However, the extent to which this objective is achieved relies on the underlying assumption that the pre-disaster variation in each dimension for each disaster conforms to a normal distribution. We conduct Shapiro–Wilk tests for the pre-disaster variation of each dimension across all disasters (24 dimensions \times 203 disasters) to verify how broadly this assumption holds the dataset. As seen in Fig. 7, we find that at a significance level of 5%, we fail to reject the null hypothesis of normality for at least 75% of the events for 23 of the dimensions. In the case of the remaining dimension, *family*, we fail to reject the normality hypothesis for pre-disaster variation in only 64% of the events. We also visually inspected Q–Q plots for a number of dimension-disaster pairs to ensure that our intuition aligned with the p value of the statistical test.

A.2: Clustering quality measurements for different numbers of clusters (N)

See Fig. 8

Fig. 8 Silhouette coefficient, Dunn Index, and Connectivity for $N \in 1 : 10$. Higher Silhouette coefficient and Dunn index values, as well as lower Connectivity, indicate higher quality clustering

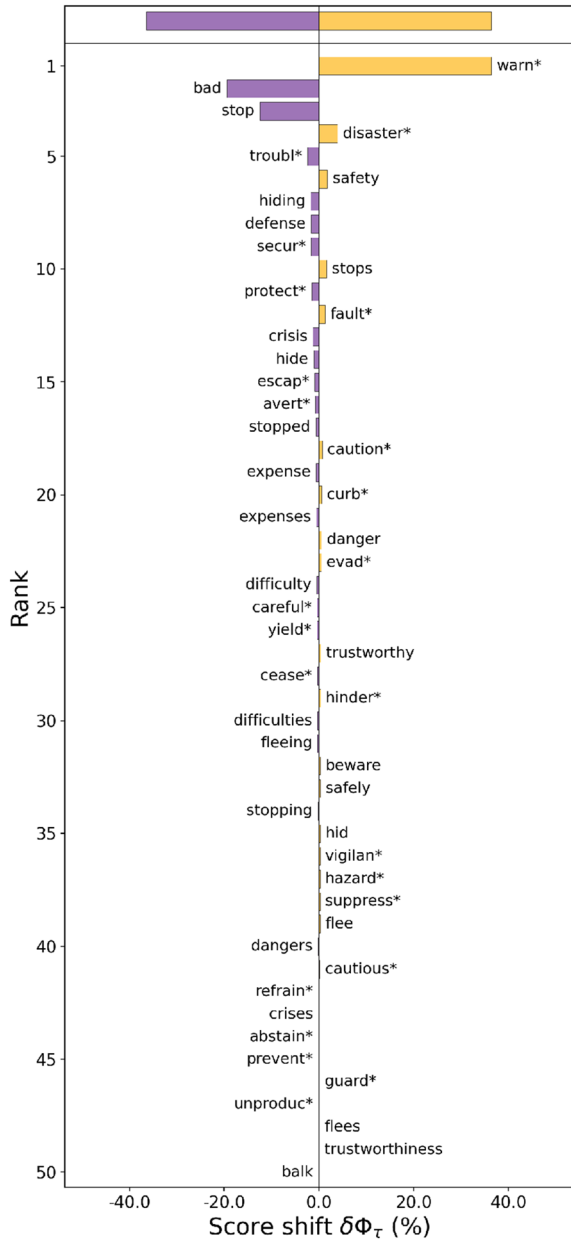


A.3: Word-shift analysis of prototypical dimensions before and after the onset of disaster

A.3.1: Risk

See Fig. 9.


Fig. 9 Top 50 overall proportional word-shifts for **Risk**



A.3.2: Achieve

See Fig. 10.


Fig. 10 Top 50 overall proportional word-shifts for **Achieve**



A.3.3: Reward

See Fig. 11.


Fig. 11 Top 50 overall proportional word-shifts for **Reward**



A.3.4: Social

See Fig. 12.


Fig. 12 Top 50 overall proportional word-shifts for **Social**



A.3.5: Health

See Fig. 13.

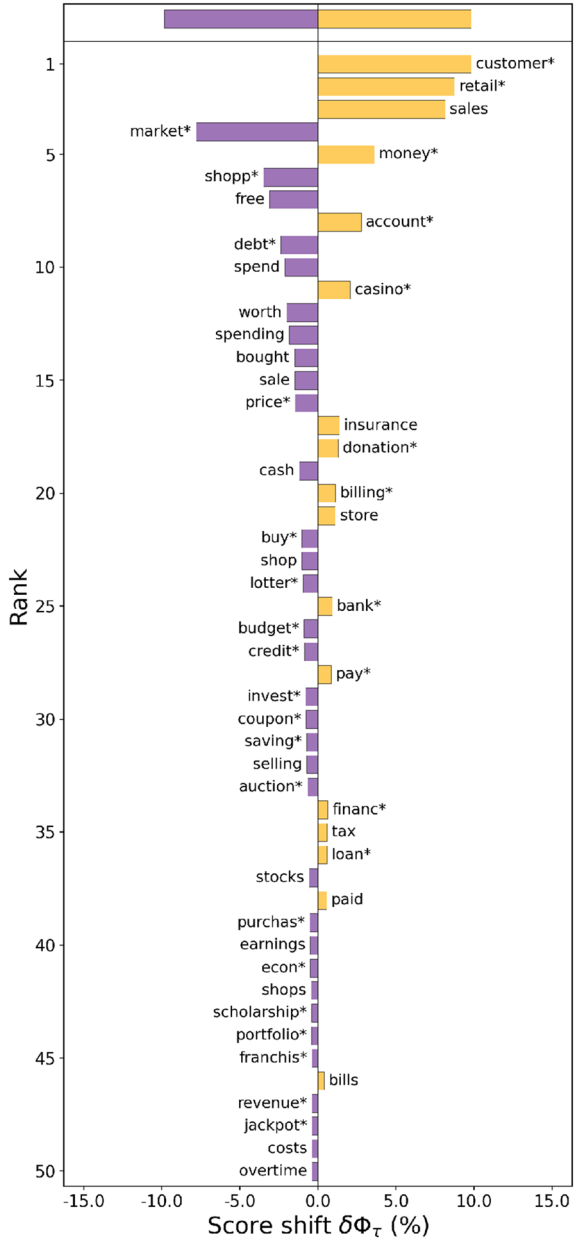
Fig. 13 Top 50 overall proportional word-shifts for **Health**



A.3.6: Money

See Fig. 14.

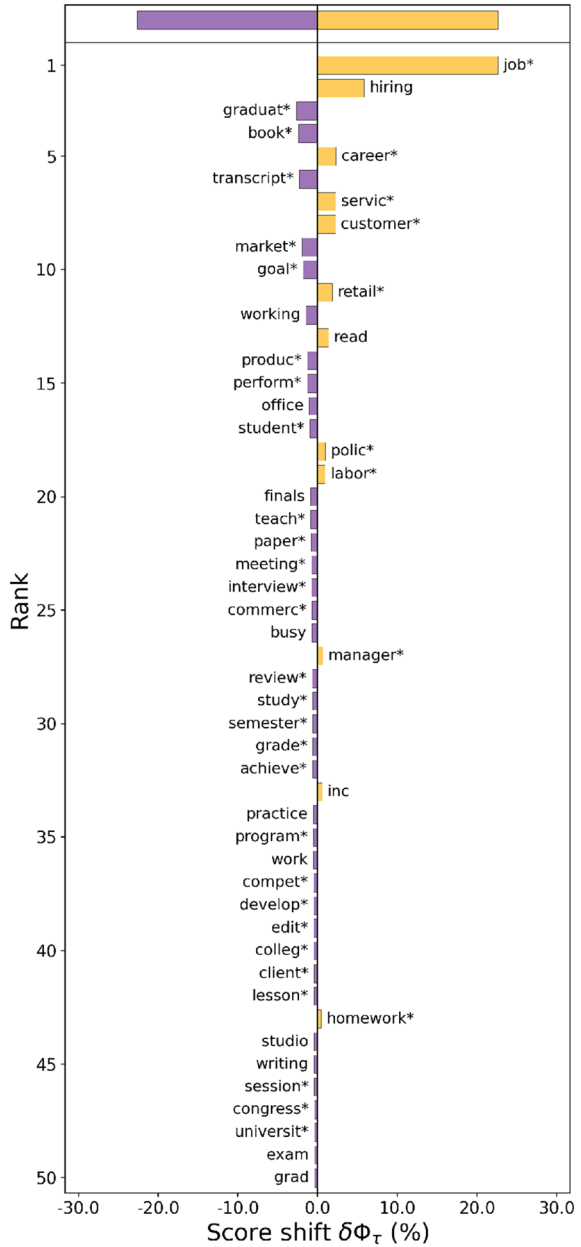
Fig. 14 Top 50 overall proportional word-shifts for **Money**



A.3.7: Work

See Fig. 15.

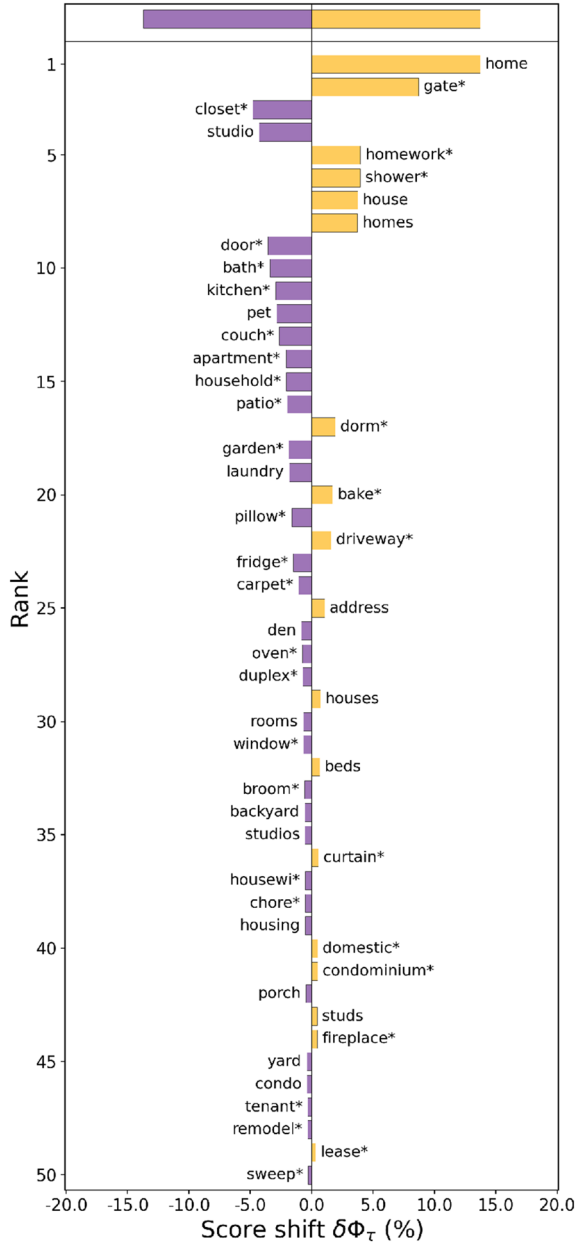
Fig. 15 Top 50 overall proportional word-shifts for **Work**



A.3.8: Home

See Fig. 16.


Fig. 16 Top 50 overall proportional word-shifts for **Home**



A.3.9: Ingest

See Fig. 17.

Fig. 17 Top 50 overall proportional word-shifts for **Ingest**



A.4: Word-shift analysis for dimensions that differ between clusters

A.4.1: Sadness

See Fig. 18.




Fig. 18 Proportional word-shifts for **Sad** for the two clusters of disasters

A.4.2: Anger

See Fig. 19.




Fig. 19 Proportional word-shifts for Anger for the two clusters of disasters

A.4.3: Swear

We have excluded the word-shifts for *swear* as they contain words that are offensive, particularly when read out of context. The observations from these word-shifts bear a close similarity to those of *anger*

A.4.4: Fear

See Fig. 20.




Fig. 20 Proportional word-shifts for **Fear** for the two clusters of disasters

A.5: Exemplar tweets for selected lexicon dimensions

The text of the tweets shown below have been anonymized by removing usernames and identifying hashtags as well as by paraphrasing where necessary.

A.5.1: We

1. RT <username>: What are some inspiring scriptures about HOPE?? We all need a little bit of hope today
2. Tornado sirens. Here we go I guess
3. <username> I posted this once we got back to the dorms, crazy
4. <username> Hold on! We need you all to give the rest of us hope and bring some sense to the state
5. RT <username>: We can decide to be hopeful—don't have to wait until we feel like it

A.5.2: Ingest

1. He is just feeling upset because he didn't get free *food* from red cross
2. First food in 3 days. (at <username>)
3. I can't drive in this how tf am I supposed to eat!!!!?
4. Dinner is served! #paleo
5. I'm hungry.. ima cry

A.5.3: Health

1. ... just finished pushing my workout to the next level with <hashtag>
2. Appointment to see the dentist today....oh joy!
3. ... gets my mind right. <hashtag> <hashtag> <hashtag> <hashtag> #workout
4. I've never been so tired
5. Smooth ride to... first doctor visit!

A.5.4: Sad in Cluster 1

1. Tonight was seriously so depressing.
2. I can't hold the tears back anymore...
3. <username> this makes me so sad
4. ... yesterday was a completely strange & weird day < hashtag> #<city>flood <hashtag> <hashtag> #sad
5. Lost my <hashtag>... poster from <hashtag> in last night's flooding. Had that thing for... years.):

A.5.5: Sad in Cluster 2

1. Feel so lonely... I miss my... cause they're there when I need them <sad emoji>

2. This is so sad they have to save people before saving animals # <city>flood
3. Sigh I miss the sunny & hot <city>.
4. This is the worst news I can wake up to. I'll miss you <name>.
5. So sad to hear about the... lost their lives during... tornado in <state>. What a tragedy

A.5.6: Anger in Cluster 1

1. Monday.. means I have school *shit*... my day is ruined
2. Seriously don't tell me *shit* like that and then ignore me
3. so many people give him *shit*, just because he's deaf. like what the fuck?
4. <username>: First, no one left early. I was there. Second, stop talking about this *shit*
5. buses are always up in my *shit* until I need to catch one

A.5.7: Anger in Cluster 2

1. People in <state> only give a *fuck* about others when disasters happen.
2. <username> it's *pissing* me the *fuck* off!!
3. Ugh *hate* feeling helpless
4. It's *frustrating* when people try to politicize a horrific disaster
5. I *hate* the way <state> handles disasters #Unorganized. Lack of communication makes my job easier not harder


A.5.8: Anxiety in Cluster 1

1. I've never been so *scared* in my life from thunder.
2. <username> I'm starting to get *scared* now
3. RT <username> Most attention being paid to Sunday's heavy rain *threat*, but locally heavy rain is possible today too...
4. I was so depressed and upset yesterday, I went to bed by 5:30 and slept until 4am. Woke up still angry
5. Overwhelmed trying to figure out how I'll catch morning flight w/o transit or shuttle to airport. Here's hoping service resumes.

A.5.9: Anxiety in Cluster 2

1. <username> Yeah, my blood *pressure* is up
2. Days like this can really raise my *pressure*
3. Idk what's going on but I am lost and *confused*:(#badday
4. Really *upset* right now. I miss <country>. I miss warmer weather not shitty rain and the cold. I miss my family!
5. News showing a tornado hit right where my boyfriend is at work. About to have a *panic* attack.:(

Fig. 21 Comparison of forecast performance between the predicted cluster trajectory and the baseline forecasting method. The heights of the bars indicate the percentage of times each method provided a better forecast of community response to a particular event for different lengths of observed segments



A.5.10: Body in Cluster 1

1. While you were *asleep*...#<state>flood
2. Thank you flash flood warning for giving me a *heart* attack as I am trying to fall *asleep*! I am now wide awake
3. I am so excited about getting to sleep in MY OWN BED tomorrow
4. Its not incense, its a heating pad filled with seeds and lavender to help sore *muscles* and headaches
5. This chilly weather and painful ribs just make me want to crawl in bed and snuggle all afternoon

A.5.11: Body in Cluster 2

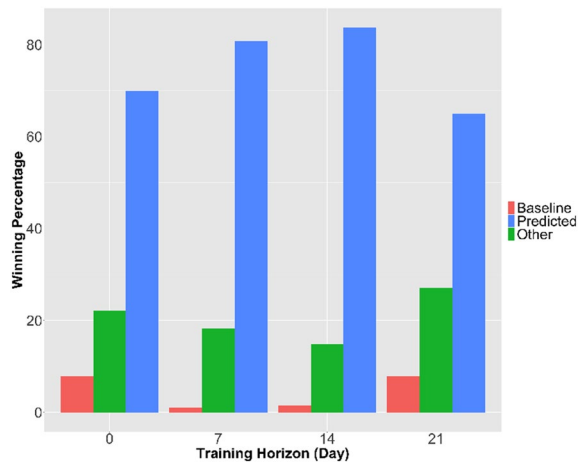
1. I've had one shower and only 8hrs *sleep* maybe in 3 days
2. Gooood news! My sunburn is gone and I'm not *peeling*. perks of flooding your *body* with water and 18+ layers of lotion applied!
3. house flooded and alone at a hotel? take selfies with your bare *face* <url>
4. RT <user>: My *heart* goes out to everyone impacted by tornadoes in the <city> area. Stay safe and #PrayFor<state>.
5. Don't think I'll be going to sleep anytime soon

A.6: Forecasting community response using cluster trajectories

We adapt an approach inspired by cross-validation for statistical models to evaluate the usefulness of the clustered community response trajectories in forecasting the community response for future events.

We repeated the time series clustering N ($=203$) times, each time leaving one event out. As with the original results discussed in the paper, each time we used hierarchical clustering with DTW distance and set the number of clusters to 2. Then,

Fig. 22 Three-way comparison of forecasting performance across the predicted cluster trajectory, the alternative cluster trajectory (i.e., the one that was not chosen), and the baseline forecasting method. The heights of the bars indicate the percentage of times each method provided a better forecast of the community response to a particular event for different lengths of observed segments



for each clustering result, we estimate the DTW distance between different segments of the unobserved response that begin from the start of the study period and the representative signatures of the two clusters. Specifically, we considered three segments of length: 1, 2, and 3 weeks (i.e., $0 \leq D < 7$, $0 \leq D < 14$, $\leq D < 21$). We select the cluster that has the smaller DTW distance to the observed segment of the test event as our prediction of how the community response of that event would have evolved in the unobserved period.

For each observed segment length, we compare the similarity of the unobserved segment of community response to the selected cluster, the remaining cluster, and a baseline prediction. The baseline prediction was generated by individually fitting ARIMA models to the different dimensions (23) of the observed segment.⁷ As Fig. 21 shows, the response cluster that was most similar to the observed segment of the test event is a better predictor of how the community response evolved from there during the unobserved period consistently for more than 75% of the events across different lengths of the observed period. The performance of the cluster prediction relative to the baseline model is best when a moderate amount of the community response has been observed. In Fig. 21, these correspond to training horizons of Day 7 and Day 14 (i.e., lengths of the observed period are 2 and 3 weeks, respectively). This suggests that in the early stages of a community response to disaster, cluster assignment benefits more from incremental information. The comparative decline in cluster-based predictive performance at the training horizon of Day 21 (i.e., a 4-week observed period, and predictions are made for only 7 days) is likely to be due to the baseline model continuing to be fit to the new data, whereas, on average, the influence of those data on cluster assignment decreases. Also, the cluster trajectories themselves are fixed and do not benefit from additional information. It remains an open question as to whether the baseline model would do better or worse than it is now (winning less than 25% of the one-versus-one comparisons against cluster

⁷ We used the **ARIMA** function from the R package *fable*.

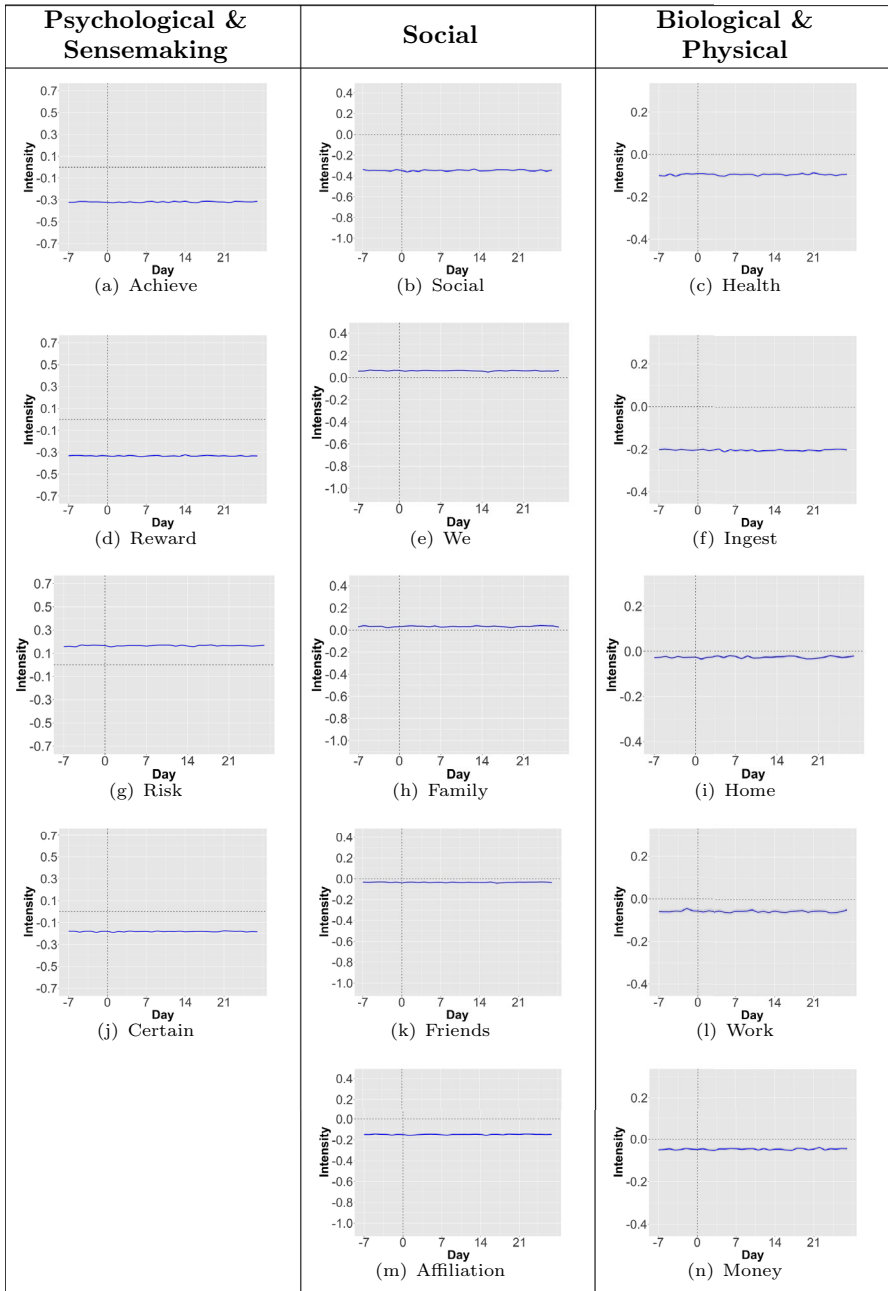


Fig. 23 Community response patterns using shuffled observations for dimensions with universal trajectories across disasters in the original data. All figures in a column have the same intensity (Y) scales which are identical to those of Fig. 5. The shaded areas are 95% bootstrap confidence intervals. If a confidence interval is not visible for a particular dimension, it is narrower than width of the line showing the mean trajectory

prediction), if community response is forecasted for more than 1 week while holding the length of the observed period at 4 weeks. In addition to comparing the forecasting performance of the assigned cluster against the baseline model, we also conduct a three-way comparison between both clusters and the baseline model. Figure 22 shows that the baseline model performs even worse than before when compared with the forecasts of both clusters. In other words, even when the baseline model provides a better forecast than the prediction of the assigned cluster, it is frequently beaten by the forecast offered by the cluster that was not chosen. This suggests that the two cluster-based predictions, when taken together, as a primary forecast and an alternate forecast, may be more useful than only the forecast from the assigned cluster.

A.7: Validation of observed community response time series

Among the results of this study, the lexicon time series shown in Fig. 5 constitute community response patterns that generalize well across the disasters in our data. In this section, we validate that these patterns are not artifacts of data generation or processing by comparing them against the corresponding time series generated from a null model.

We generate a new time series for each dimension of each time series by shuffling the daily observations. We only shuffle observations within the study period once they have been normalized w.r.t. the behavior from the 3-week baseline period). We repeat this procedure 200 times across all disasters to generate as many sets of community responses. Then, we use each synthesized set of community responses to generate representative time series for each of the dimensions in Fig. 5. The result is a distribution of 200 synthetic representative time series for each dimension. In Fig. 23, we replicate Fig. 5 using the synthetic distributions. The figure shows that shuffled observations result in synthetic community responses that, for the most part, do not vary with time and have none of the rich temporal structure observed in the actual data. For all dimensions, the community responses correspond to a fixed intensity, which corresponds to the average level in relation to the baseline period.

Funding This work is supported in part by the Air Force Office of Scientific Research under Award No. FA9550-19-1-0029.

Data Availability The tweet data generated during the current study are not publicly available owing to privacy issues and Twitter terms of service. However, corresponding tweet identifiers are available from the corresponding author upon reasonable request.

Declarations

Conflict of interest On behalf of all authors, the corresponding author states that there is no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Bonanno, G. A., Brewin, C. R., Kaniasty, K., & Greca, A. M. L. (2010). Weighing the costs of disaster: Consequences, risks, and resilience in individuals, families, and communities. *Psychological Science in the Public Interest*, *11*(1), 1–49.
2. Norris, F. H., Friedman, M. J., Watson, P. J., Byrne, C. M., Diaz, E., & Kaniasty, K. (2002). 60,000 disaster victims speak: Part i. an empirical review of the empirical literature, 1981–2001. *Psychiatry: Interpersonal and Biological Processes*, *65*(3), 207–239.
3. Tierney, K. J., Lindell, M. K., & Perry, R. W. (Eds.). (2001). *Facing the unexpected: Disaster preparedness and response in the United States*. The National Academies Press.
4. Council, N. R. (2006). *Facing hazards and disasters: Understanding human dimensions*. The National Academies Press.
5. Perry, R. W. (2018). Defining disaster: An evolving concept. In: H. Rodríguez, W. Donner, & J. E. Trainor (Eds.), *Handbook of disaster research* (pp. 3–22). Springer.
6. DeWolfe, D. J., Nordboe, D., & Center for Mental Health Services (U.S.). (2000). *Training manual for mental health and human service workers in major disasters*. 2nd Ed. Rockville, Md.: U.S. Dept. of Health and Human Services, Substance Abuse and Mental Health Services Administration, Center for Mental Health Services.
7. Dynes, R. R. (1970). *Organized behavior in disaster*. Disaster Research Center series. Lexington, Mass: Heath Lexington Books.
8. Alexander, D. (1993). *Natural disasters*. London: UCL Press.
9. Kates, R. W., & Pijawka, D. (1977). From rubble to monument: The pace of reconstruction. *Reconstruction Following Disaster*, *1*, 1–23.
10. Bonanno, G. A. (2004). Loss, trauma, and human resilience: Have we underestimated the human capacity to thrive after extremely aversive events? *American Psychologist*, *59*(1), 20.
11. Kogan, M., Palen, L., & Anderson, K. M. (2015). Think local, retweet global: Retweeting by the geographically-vulnerable during hurricane sandy. In: *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, pp. 981–993. ACM.
12. Houston, J. B., Hawthorne, J., Perreault, M. F., Park, E. H., Goldstein Hode, M., Halliwell, M. R., Turner McGowen, S. E., Davis, R., Vaid, S., McElderry, J. A., et al. (2015). Social media and disasters: A functional framework for social media use in disaster planning, response, and research. *Disasters*, *39*(1), 1–22.
13. Lin, Y.-R., Margolin, D., & Wen, X. (2017). Tracking and analyzing individual distress following terrorist attacks using social media streams. *Risk Analysis*, *37*(8), 1580–1605.
14. Brubaker, J., Kivran-Swaine, F., Taber, L., & Hayes, G. (2012). Grief-stricken in a crowd: The language of bereavement and distress in social media. In: *Proceedings of the International AAAI Conference on Web and Social Media*, *6*, 42–49.
15. Lin, Y.-R., & Margolin, D. (2014). The ripple of fear, sympathy and solidarity during the Boston bombings. *EPJ Data Science*, *3*(1), 31.
16. Olteanu, A., Vieweg, S., & Castillo, C. (2015). What to expect when the unexpected happens: Social media communications across crises. In: *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, pp. 994–1009. ACM.

17. Forbes, D., Alkemade, N., Waters, E., Gibbs, L., Gallagher, C., Pattison, P., Lusher, D., MacDougall, C., Harms, L., Block, K., et al. (2015). The role of anger and ongoing stressors in mental health following a natural disaster. *Australian & New Zealand Journal of Psychiatry*, 49(8), 706–713.
18. Grimm, A., Hulse, L., & Schmidt, S. (2012). Human responses to disasters: A pilot study on peri-traumatic emotional and cognitive processing. *Europe's Journal of Psychology*, 8(1), 112.
19. Anderson, T. J., Kogan, M., Bica, M., Palen, L., Anderson, K. M., Morss, R., Demuth, J., Lazrus, H., Wilhelmi, O., & Henderson, J. (2016). Far far away in far rockaway: Responses to risks and impacts during hurricane sandy through first-person social media narratives. In: *ISCRAM*.
20. Yule, W., Bolton, D., Udwin, O., Boyle, S., O'Ryan, D., & Nurrish, J. (2000). The long-term psychological effects of a disaster experienced in adolescence: I: The incidence and course of ptsd. *Journal of Child Psychology and Psychiatry*, 41(4), 503–511.
21. Cherry, K. E., Calamia, M. R., Birch, T., & Moles, A. (2020). *Assessing mental health after a disaster: Flood exposure, recovery stressors, and prior flood experience* (pp. 271–283). Springer International Publishing. https://doi.org/10.1007/978-3-030-51525-6_16
22. Kar, N., & Bastia, B. K. (2006). Post-traumatic stress disorder, depression and generalised anxiety disorder in adolescents after a natural disaster: A study of comorbidity. *Clinical Practice and Epidemiology in Mental Health*, 2(1), 1–7.
23. (2014). *DSM-5 handbook of differential diagnosis*. Washington, D.C.: American Psychiatric Publishing.
24. Drabek, T. E. (2018). Community processes: Coordination. In: H. Rodriguez, W. Donner, J. Trainor (Eds.), *Handbook of disaster research. Handbooks of sociology and social research*. Springer, Cham. https://doi.org/10.1007/978-3-319-63254-4_25
25. Horlick-Jones, T. (1995). Modern disasters as outrage and betrayal. *International Journal of Mass Emergencies and Disasters*, 13(3), 305–315.
26. Bourque, L. B., Siegel, J. M., Kano, M., & Wood, M. M. (2006). Weathering the storm: The impact of hurricanes on physical and mental health. *The Annals of the American Academy of Political and Social Science*, 604(1), 129–151.
27. Lindell, M. K., & Perry, R. W. (2004). *Communicating environmental risk in multiethnic communities* (Vol. 7). SAGE Publications. <https://doi.org/10.4135/9781452229188>
28. van den Berg, B., Grievink, L., Yzermans, J., & Leuret, E. (2005). Medically unexplained physical symptoms in the aftermath of disasters. *Epidemiologic Reviews*, 27(1), 92–106.
29. Lindell, M. K. (2013). Recovery and reconstruction after disaster. In: P. T. Bobrowsky (Eds.) *Encyclopedia of natural hazards*. Encyclopedia of Earth Sciences Series. Springer, Dordrecht. https://doi.org/10.1007/978-1-4020-4399-4_285
30. Maslow, A. H. (1954). *Motivation and personality*. Harper's psychological series. New York: Harper & Row.
31. Townshend, I., Awosoga, O., Kulig, J., & Fan, H. (2015). Social cohesion and resilience across communities that have experienced a disaster. *Natural Hazards*, 76(2), 913–938.
32. Myers, D. G. (1994). *Disaster response and recovery: A Handbook for Mental Health Professionals*. DIANE Publishing.
33. Raphael, B. (1986). *When disaster strikes: how individuals and communities cope with catastrophe*. New York: Basic Books.
34. Williams, D. (1999). Human responses to change. *Futures*, 31(6), 609–616.
35. Sutton, J., Gibson, C. B., Phillips, N. E., Spiro, E. S., League, C., Johnson, B., Fitzhugh, S. M., & Butts, C. T. (2015). A cross-hazard analysis of terse message retransmission on twitter. *Proceedings of the National Academy of Sciences*, 112(48), 14793–14798.
36. Faas, A. J., & Jones, E. C. (2017). Chapter 2—Social network analysis focused on individuals facing hazards and disasters. In E. C. Jones & A. J. Faas (Eds.), *Social network analysis of disaster response, recovery, and adaptation* (pp. 11–23). Butterworth-Heinemann. <https://doi.org/10.1016/B978-0-12-805196-2.00002-9>.
37. Shklovski, I., Palen, L., & Sutton, J. (2008). Finding community through information and communication technology in disaster response. In: *Proceedings of the 2008 ACM Conference on Computer Supported Cooperative Work*, pp. 127–136. ACM.
38. Starbird, K., & Palen, L. (2010). Pass it on?: Retweeting in mass emergency. In: *7th International Conference on Information Systems for Crisis Response and Management: Defining Crisis Management 3.0, Proceedings*. ISCRAM.
39. Lobb, A., Mock, N., & Hutchinson, P. L. (2012). Traditional and social media coverage and charitable giving following the 2010 earthquake in Haiti. *Prehospital and Disaster Medicine*, 27(04), 319–324.

40. Imran, M., Castillo, C., Diaz, F., & Vieweg, S. (2015). Processing social media messages in mass emergency: A survey. *ACM Computing Surveys (CSUR)*, 47(4), 67.
41. Imran, M., Elbassuoni, S., Castillo, C., Diaz, F., & Meier, P. (2013). Practical extraction of disaster-relevant information from social media. In: *Proceedings of the 22nd International Conference on World Wide Web*, pp. 1021–1024. ACM.
42. Olteanu, A., Castillo, C., Diaz, F., & Vieweg, S. (2014). Crisislex: A lexicon for collecting and filtering microblogged communications in crises. In: *ICWSM*.
43. Murthy, D., & Longwell, S. A. (2013). Twitter and disasters: The uses of twitter during the 2010 Pakistan floods. *Information, Communication & Society*, 16(6), 837–855.
44. David, C. C., Ong, J. C., & Legara, E. F. T. (2016). Tweeting Supertyphoon Haiyan: Evolving functions of twitter during and after a disaster event. *PLoS One*, 11(3), 0150190.
45. De Choudhury, M., Monroy-Hernandez, A., & Mark, G. (2014). Narco emotions: affect and desensitization in social media during the Mexican drug war. In: *Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems*, pp. 3563–3572. ACM.
46. Wen, X., & Lin, Y.-R. (2016). Sensing distress following a terrorist event. In: *International Conference on Social Computing, Behavioral-Cultural Modeling and Prediction and Behavior Representation in Modeling and Simulation*, pp. 377–388. Springer.
47. Garcia, D., & Rimé, B. (2019). Collective emotions and social resilience in the digital traces after a terrorist attack. *Psychological Science*, 30(4), 617–628.
48. Saha, K., & De Choudhury, M. (2017). Modeling stress with social media around incidents of gun violence on college campuses. *Proceedings of the ACM on Human-Computer Interaction*, 1(CSCW), 92.
49. De Choudhury, M., & Counts, S. (2013). Understanding affect in the workplace via social media. In: *Proceedings of the 2013 Conference on Computer Supported Cooperative Work*, pp. 303–316. ACM.
50. Cohn, M. A., Mehl, M. R., & Pennebaker, J. W. (2004). Linguistic markers of psychological change surrounding September 11, 2001. *Psychological Science*, 15(10), 687–693.
51. Pellert, M., Metzler, H., Matzenberger, M., & Garcia, D. (2022). Validating daily social media macroscopes of emotions. *Scientific Reports*, 12(1), 11236.
52. Garcia, D., Pellert, M., Lasser, J., & Metzler, H. (2021). Social media emotion macroscopes reflect emotional experiences in society at large. arXiv preprint [arXiv:2107.13236](https://arxiv.org/abs/2107.13236).
53. Mohammad, S. M., & Turney, P. D. (2013). Crowdsourcing a word-emotion association lexicon. *Computational Intelligence*, 29(3), 436–465.
54. Norris, F. H., Phifer, J. F., & Kaniasty, K. (1994). Individual and community reactions to the Kentucky floods: findings from a longitudinal study of older adults. In *Individual and community responses to trauma and disaster: The structure of human chaos* (pp. 378–400). Cambridge University Press. <https://doi.org/10.1017/CBO9780511570162.019>.
55. Grimm, A., Hulse, L., Preiss, M., & Schmidt, S. (2014). Behavioural, emotional, and cognitive responses in European disasters: results of survivor interviews. *Disasters*, 38(1), 62–83.
56. Gortner, E.-M., & Pennebaker, J. W. (2003). The archival anatomy of a disaster: Media coverage and community-wide health effects of the Texas a & m bonfire tragedy. *Journal of Social and Clinical Psychology*, 22(5), 580–603.
57. Richardson, B. K., Siebeneck, L. K., Shaunfield, S., & Kaszynski, E. (2014). From no man's land to a stronger community: Communitas as a theoretical framework for successful disaster recovery. *International Journal of Mass Emergencies & Disasters*, 32(1), 194–219.
58. Coffelt, T. A., Smith, F. L., Sollitto, M., & Payne, A. R. (2010). Using sensemaking to understand victims' responses to a natural disaster. *Northwest Journal of Communication*, 39(1), 11–35.
59. Kuijter, R. G., & Boyce, J. A. (2012). Emotional eating and its effect on eating behaviour after a natural disaster. *Appetite*, 58(3), 936–939.
60. Kotsakos, D., Trajcevski, G., Gunopulos, D., & Aggarwal, C. C. (2018). Time-series data clustering. In *Data Clustering* (pp. 357–380). Chapman and Hall/CRC.
61. Petitjean, F., Ketterlin, A., & Gançarski, P. (2011). A global averaging method for dynamic time warping, with applications to clustering. *Pattern Recognition*, 44(3), 678–693.
62. Satter, E. (2007). Hierarchy of food needs. *Journal of Nutrition Education and Behavior*, 39(5), 187–188.
63. Dobie, S., Schneider, J., & Szafranski, A. (2019). Going beyond the waffle house index: Using food systems as an indicator of community health and sustainability. In: *2019 IEEE International Symposium on Technologies for Homeland Security (HST)*, pp. 1–6. IEEE.
64. Kohn, R., & Levav, I. (1990). Bereavement in disaster: An overview of the research. *International Journal of Mental Health*, 19(2), 61–76.

65. Gaspar, R., Pedro, C., Panagiotopoulos, P., & Seibt, B. (2016). Beyond positive or negative: Qualitative sentiment analysis of social media reactions to unexpected stressful events. *Computers in Human Behavior*, *56*, 179–191.
66. Saleem, H. M., Xu, Y., & Ruths, D. (2014). Effects of disaster characteristics on twitter event signature. *Procedia Engineering*, *78*, 165–172.
67. Sloan, L., & Morgan, J. (2015). Who tweets with their location? understanding the relationship between demographic characteristics and the use of geoservices and geotagging on twitter. *PLoS One*, *10*(11), 0142209.
68. Walker, M., & Matsa, K. E. (2021). News consumption across social media in 2021.
69. Hagon, K., Turmine, V., Pizzini, G., Freebairn, A., & Singh, R. (2020). Hazards everywhere – climate and disaster trends and impacts. In: *World Disasters Report: Come Heat or High water*. International Federation of Red Cross, Geneva. Chap. 2. OCLC: 1232500224.
70. Ritchie, H., & Roser, M. (2014). Natural disasters. Our world in data. <https://ourworldindata.org/natural-disasters>. Accessed 30 Jan 2023
71. Smith, A. B. (2020). US Billion-Dollar Weather and Climate Disasters, 1980–Present (NCEI Accession 0209268). NOAA National Centers for Environmental Information, 10.
72. Coronese, M., Lamperti, F., Keller, K., Chiaromonte, F., & Roventini, A. (2019). Evidence for sharp increase in the economic damages of extreme natural disasters. *Proceedings of the National Academy of Sciences*, *116*(43), 21450–21455.
73. Nietzsche, F., & Levy, O. (1909). *The twilight of the idols* (Vol. 17). Allen & Unwin.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Authors and Affiliations

Danaja Maldeniya¹  · Munmun De Choudhury⁴ · David Garcia^{5,6,7} · Daniel M. Romero^{1,2,3}

✉ Danaja Maldeniya
dmal@umich.edu

Munmun De Choudhury
mchoudhu@cc.gatech.edu

David Garcia
garcia@csh.ac.at

Daniel M. Romero
drom@umich.edu

¹ School of Information, University of Michigan, Ann Arbor, MI, USA

² Center for the Study of Complex Systems, University of Michigan, Ann Arbor, MI, USA

³ Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI, USA

⁴ School of Interactive Computing, Georgia Tech, Atlanta, GA, USA

⁵ Complexity Science Hub, Vienna, Austria

⁶ Medical University of Vienna, Vienna, Austria

⁷ Graz University of Technology, Vienna, Austria