Communicating Hurricane Risks: Multi-Method Examination of Risk Imagery Diffusion

Melissa Bica
University of Colorado Boulder
Boulder, CO, USA
melissa.bica@colorado.edu

Julie L. Demuth
National Center for Atmospheric Research
Boulder, CO, USA
jdemuth@ucar.edu

James E. Dykes
University of Colorado Boulder
Boulder, CO, USA
james.dykes@colorado.edu

Leysia Palen
University of Colorado Boulder
Boulder, CO, USA
palen@colorado.edu

ABSTRACT

Conveying uncertainty in information artifacts is difficult; the challenge only grows as the demand for mass communication through multiple channels expands. In particular, as natural hazards increase with changing global conditions, including hurricanes which threaten coastal areas, we need better means of communicating uncertainty around risks that empower people to make good decisions. We examine how people share and respond to a range of visual representations of risk from authoritative sources during hurricane events. Because these images are now shared widely on social media platforms, Twitter provides the means to study them on a large scale as close to *in vivo* as possible. Using mixed methods, this study analyzes diffusion of and reactions to forecast and other risk imagery during the highly damaging 2017 Atlantic hurricane season to describe the collective response to visual representations of risk.

CCS CONCEPTS

• **Human-centered computing → Empirical studies in collaborative and social computing**;

KEYWORDS

Images, Information Diffusion, Risk Communication

CHI 2019, May 4–9, 2019, Glasgow, Scotland UK

ACM ISBN 978-1-4503-5970-2/19/05…$15.00
https://doi.org/10.1145/3290605.3300545

ACM Reference Format:
Melissa Bica, Julie L. Demuth, James E. Dykes, and Leysia Palen. 2019. Communicating Hurricane Risks:, Multi-Method Examination of Risk Imagery Diffusion. In *CHI Conference on Human Factors in Computing Systems Proceedings (CHI 2019), May 4–9, 2019, Glasgow, Scotland UK*. ACM, New York, NY, USA, 13 pages. https://doi.org/10.1145/3290605.3300545

1 INTRODUCTION

Disaster events arising from natural hazards are characterized by uncertainty, such as who or where will be impacted, how and when to evacuate, or how much damage will be sustained. Risk communication therefore is critical, as its purpose is to provide information about a potential hazard and its impacts for people to use to protect themselves and mitigate destruction. However, communicating risk is itself challenging, as it often includes depicting probabilities or a range of scenarios in ways that people can understand and make decisions about. Communicating natural hazard risks has been and remains a central concern to emergency practitioners and weather scientists alike, no matter the mode of communication. Visualizations with maps and satellite imagery are typical forms of communication to reach populations in the broad swaths of geography that are threatened by hazards. Such images are “products” that are issued by formal organizations periodically and distributed to news stations and emergency management groups to use locally [13].

Hurricanes are the focus of this study because of the tremendous risk they pose to people’s well-being. Moreover, hurricane attributes such as their size, motion, intensity, and associated risks are predictable to an extent, albeit with uncertainty, yet they all can rapidly change leading up to landfall. Thus, weather forecasts and related products are distilled representations of risk communication about evolving hurricane threats [7, 8, 38]. Here we look at different types of visual hurricane risk information, as shown in Figure 1, including hurricane forecasts (e.g. the cone of uncertainty and ensemble or “spaghetti” models), observations of the hurricane (e.g. radar or satellite imagery) which inform risk

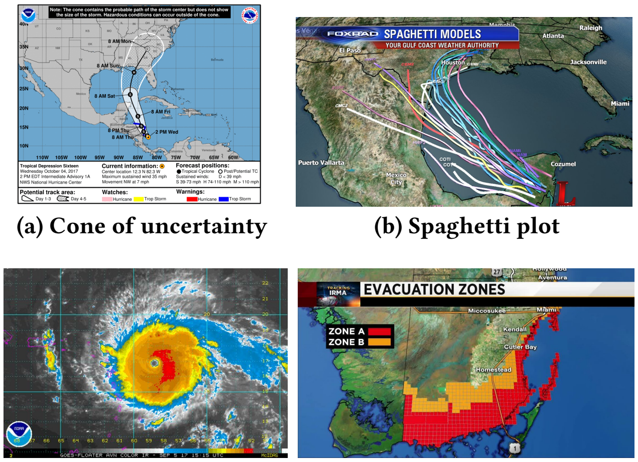

(a) Cone of uncertainty (b) Spaghetti plot

(c) Satellite image (d) Evacuation map

Figure 1: Example hurricane risk images. Credit: NOAA, FOX 26 Houston, WPLG Local 10.

assessment and forecasts, and evacuation information. We refer collectively to these as **hurricane risk images**.

People in hurricane-threatened regions are frequently exposed to these images, which have long been available via television stations, newspapers, and the National Hurricane Center (NHC) website. In recent years, social media has become another important means for distributing hurricane risk information, especially given that hurricanes typically form several days before landfall, offering much opportunity for people to discuss their risk and make meaning of these information products. Despite the prevalence and variety of hurricane risk images, the emergency management and weather science communities know little about the effectiveness of visual risk communication in supporting risk assessment and decision-making.

This multi-method investigation empirically examines the diffusion of visual risk communication products during the sequence of hurricane events during the 2017 Atlantic Hurricane season to understand which products diffuse and how. We investigate aspects of individual and collective human interaction with imagery that is generated and shared by authoritative sources, using Twitter as the platform for capturing this activity. As part of the research design, we collected tweets from these sources to focus on scientifically-informed representations. However, as we discuss in Section 7, the data collected *in response* to these posts were not restricted in any way. To begin, we inductively examined the images:

> **RQ1**: What kinds of hurricane forecast and other risk graphics are posted by authoritative sources?

Then, to understand human interaction with hurricane risk imagery, we first had to measure the diffusion of such images on Twitter. Because there are many different kinds of risks visualized for a hurricane, such as its potential track or current location, its impacts if/when it occurs, and protective and preparedness information provided to reduce one's harm from it, we also want to know how the response differs to images portraying different kinds of risk. Diffusion of risk imagery is likely to be driven by a number of factors, including but not limited to the type of risk depicted in the image and the type of authoritative source from which it comes, as there are differences in the function and mission of different groups who convey hurricane risk [13]. Thus, we ask a second set of research questions:

> **RQ2**: How do different types of authoritative-source forecast and risk images diffuse for hurricane events? How does diffusion differ based on the type of image, or the authoritative source user?

Finally, after characterizing *what* types of imagery were diffused by *whom* and *how* they diffused, we sought to understand more about *why* the images diffused as they did:

> **RQ3**: What does the content of replies and retweets/quote tweets reveal about how people relate to hurricane risk and forecast images?

We conducted three studies to answer each of the research questions: a descriptive statistical analysis of a hand-coded dataset of authoritative hurricane risk images (RQ1, Section 4), a bivariate statistical analysis of the diffusion of these images using new diffusion metrics developed for this research (RQ2, Section 5), and targeted qualitative analyses of diffused content (RQ3, Section 6).

Findings contribute to multiple concerns in the expanding field of human-computer interaction, including that of information visualization and challenges with representing concepts like uncertainty, risk, and probability in consumable ways for a broad public [19, 25]; the relationship between formal and informal risk communication by computer-mediated means in disaster response [20, 44]; understanding the nature of information diffusion on widely adopted social platforms, and devising methods for that study [17, 48]; and supporting needs of the weather scientific and practitioner communities to understand socio-behavioral phenomena around their information products [42].

2 BACKGROUND

We organize the background review section to first consider what the research reveals about the effectiveness of hurricane risk imagery in general. We then segue to studies in crisis informatics that attend to risk communication. Finally, because this research relies on methods for examining the complicated matter of social media information diffusion, we briefly review the most relevant literature in this area.

Perceptions about Hurricane Risk Images

Research on hurricane forecast image interpretation has been mostly conducted in the laboratory [43, 46, 50, 57] and via surveys [28, 37, 39]. Research has shown that people confuse deterministic and probabilistic forecasts, but prefer forecasts that explicitly express uncertainty [39, 57]. For hurricanes, though, depictions of risk and uncertainty can be difficult to effectively communicate and interpret [13, 15]. The type of graphic used affects interpretation. Forecast uncertainty visualizations can display summarized uncertainty statistics or an ensemble of multiple data values or scenarios [43]. Summary displays like the cone of uncertainty are commonly misinterpreted as showing a hurricane's size or intensity [8], while ensemble displays such as spaghetti plots which show a range of probably hurricane tracks are misused for point-based judgments rather than identification of larger patterns [43, 50]; even forecasting experts have made these errors [45]. "Indexical" images—photos of damage—may have more impact on risk perception than iconical images—evacuation maps and forecasts—because they are perceived to provide "proof" of risk [46].

Missing from these studies is the ecological validity that comes from *in vivo* examination, which we hope to account for in part in this study. This is because in addition to cognitive perceptions of forecast visuals, context and social processes also affect interpretations of risk messaging and protective decision-making behavior [32]. Prior experience with a weather phenomenon plays a significant role [12]. Similarly, evacuation decisions are influenced not only by risk information, but also by other variables like information source, household location, and vulnerabilities [23, 34].

Visual representations are more effective for communicating risk and making risk information useful for decision making than numerical representations because they reveal data patterns, hold attention, and match the qualitative ways people judge probabilities [33, 49].

Imagery & Risk Studies in Crisis Informatics

Though a majority of research on social media and disasters has focused on text communication as it supports self-organization and peer production [29, 54, 56, 60], information verification and rumor detection [35, 36, 52], localized versus international social media conversations in disasters [9, 30], and natural language processing and machine learning techniques [26, 27], a small but growing body of research examines the visual components of such communication [40, 51, 58]. We know that images not only receive more engagement on social media platforms than text-only posts [47], but also communicate more—and more complex—information. Graphics are generally better than numerical or textual representations for communicating risk [33]. The photographs and other visual images shared by those closest to the disaster have a different quality and orientation toward accuracy than those by distant onlookers [6]. Photos of damage can provide crucial information to humanitarian groups [1]. These examinations inform how we collected data and sampled and interpreted the communications between formal and informal response as well as within the informal response.

In terms of risk information, research about the Zika virus showed that social media users desired specific, contextualized information, and that local knowledge was important to inform decision-making [21]. Risk perception around the Zika epidemic was multidimensional and speculative, highlighting the need for greater engagement of the public in risk communication [20].

Diffusion via Social Media

Methods for studying the diffusion of forecast and risk imagery based on image content and source are themselves important to this work, as diffusion indicates how and to what extent images receive attention online. Information diffusion on social media is a well-studied area. Research has modeled information diffusion on online social networks like Twitter and Digg [31], including the role of strong and weak ties [5], or determining whether "cascades" or virality can be predicted based on features of the network or content itself [10]. Goel et al. [17] defined a diffusion measure, "structural virality," to distinguish between tweets that diffuse by broadcast versus virally; it uses follower networks to infer diffusion of tweets across users. In this research, we focus on *timing* rather than follower relationships because of the rapidly evolving nature of hurricanes. We are also interested in not just retweets, but also replies and quote tweets, as the latter can include context from the diffusing user such as emotional reactions, clarifying questions, and descriptions of how the information impacts them, all of which can be analyzed qualitatively for further understanding of the reasons for diffusion.

Research on crisis-related information diffusion includes how certain memes became viral during a political uprising [55], how rumors around a crisis event propagate and the trustworthiness of information [4, 16, 36], and differences in how "locals" and "globals" share information [30, 53], all of which contribute to our method and interpretation.

3 2017 ATLANTIC HURRICANE SEASON

The 2017 Atlantic hurricane season, which officially ran from June 1 to November 30,[1] was especially active and catastrophic: ten hurricanes formed, with six classified as "major," meaning Category 3 or stronger (i.e., winds greater than 110 mph) as measured by the Saffir-Simpson Hurricane Wind Scale. Extensive damage was sustained in the south-eastern US, Mexico, Central America, South America, and many island countries and territories in the Caribbean Sea. The most active portion of the season was a period of just over ten weeks and included seven hurricanes: Harvey, Irma,

[1] https://www.nhc.noaa.gov/climo/

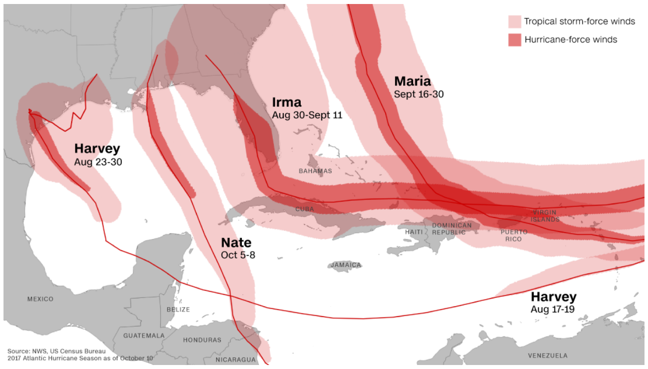

Figure 2: Tracks of Hurricanes Harvey, Irma, Maria, and Nate during the 2017 Atlantic season. Credit: CNN.

Jose, Katia, Lee, Maria, and Nate. This research focuses on four of these (see Figure 2): three which made landfall classified as major hurricanes—Harvey, Irma, and Maria—as well as Hurricane Nate, which was the costliest natural disaster ever in Costa Rica. (Jose and Lee didn't make landfall; Katia hit as a weak Category 1 with lesser damage than other events.) Thousands of fatalities occurred due to the season's hurricanes, including an estimated 2,975 alone in Puerto Rico after the massive destruction from Maria.[2]

4 STUDY 1: IMAGE COLLECTION & CODING

Data and Methods

We used a top-down approach to collect hurricane risk image tweets from accounts we identified as "authoritative sources" of information about hurricanes and their associated risks.

Identification of Authoritative Sources. Authoritative sources are people and organizations who provide authoritative and credible hurricane forecasts, observational information, and associated risks. Though others shared forecast and risk images on Twitter, we designed the scope of the investigation to examine the diffusion of and conversation around what originates as scientific communication. The primary groups who create and communicate hurricane risk and forecast information, known as the "hurricane warning system" [13] consist of the National Weather Service (NWS) forecasters at the National Hurricane Center (NHC) and local weather forecast offices who characterize and convey hurricane threats, emergency managers and other government officials who take actions (e.g., issue evacuation orders) to protect citizens, and television and other media personnel who curate and communicate hurricane information [7, 38].

We identified these sources in two ways: manually, employing the expertise of one of our authors and other collaborators who work in meteorology and weather risk communication at the National Center for Atmospheric Research

[2]https://www.bbc.com/news/world-us-canada-45338080

Table 1: Counts of Twitter accounts and risk image tweets for each authoritative source user category.

User Category	Accounts		Original Tweets	
Weather-News	188	38.4%	9433	57.1%
Weather-Gov't	41	8.4%	2890	17.5%
Weather-Other	14	2.9%	472	2.9%
News/Media	142	29.0%	3067	18.6%
Government	98	20.0%	657	4.0%
Other	6	1.2%	12	0.1%
Total	489	100%	16,531	100%

(NCAR), and via user-created lists that compile Twitter accounts of official information. For manual identification, we created lists of users for each major and/or landfalling hurricane as it occurred (Harvey, Irma, Maria, and Nate), as well as kept a running list of "general" sources that applied to all events (e.g. NWS). Though we did not identify sources specifically for the minor and/or non-landfalling hurricanes that occurred in between (Jose, Katia, and Lee), these hurricanes were included in the tweet data of many of the sources we identified. For the Twitter lists, we found two public lists each for Harvey and Irma,[3] created by twitterers @mattdpearce, a national correspondent for the Los Angeles Times, and @FEMAlive, the official FEMA account for live chats and events, which is only active during specific events.

In total, the research identified 796 Twitter accounts, classified into one of five categories: **Weather-News/Media**: TV meteorologists; websites or TV channels devoted to weather, or the weather division of broadcast news channels; **Weather-Government**: weather-related government agencies like NOAA, NASA; meteorologists/scientists at NWS, including NHC; **Weather-Other**: independent weather experts, researchers, and stormchasers; **News** (non-weather-specific): local through international news and media agencies, reporters, journalists, and photographers; TV and/or online; or **Government** (non-weather-specific): public officials; emergency management; government agencies and departments; military; city, county, and state accounts. Counts of accounts and tweets that constituted the final data set for each of these authoritative source user categories are shown in Table 1.

Tweet Data Collection. For each authoritative source, we next collected: all tweets *from* the user, all tweets *in reply to* the user, and all *retweets* and *quote tweets* (similar to retweets, but including extra comments from the reposter) of the user's tweets—what we collectively refer to as a user's *contextual-plus tweet stream*. The "plus" refers to the extension of the data collection method for "contextual tweet streams," which

[3]Hurricanes Maria and Nate did not have similar Twitter lists, and so we manually identified local sources of risk information for these.

entails collecting all tweets posted *from* the set of identified users over some window of time [3]. By additionally collecting the replies to and retweets of these authoritative users' tweets, we can study their diffusion quantitatively and qualitatively, including the conversations that form around users' original tweets. Contextual-plus and simple contextual tweet streams identify a far larger proportion of event-related tweets than relying on keywords and hashtags alone [59], and they typically contain more contextual information about a user's experience with the event [2].

We collected the contextual-plus tweet streams of each of the 796 authoritative starting when Harvey first formed to when Nate dissipated (Aug. 17–Oct. 10, 2017), totaling 9,866,351 tweets. We are interested in risk information and responses pertaining to the forecast phase of a hurricane—i.e., when the hurricane is threatening and occurring—rather than the post-disaster response and recovery phases. So, though the impacts of each hurricane continued long after the final date chosen for this research, the content informing risk assessment and preparatory decision-making was produced in the time leading up to landfall.

Coding Images for Hurricane Forecast/Risk Information. The first round of coding identified all tweets posted by authoritative sources containing hurricane risk imagery (either as still images, videos, or animated GIFs). Of the 9.87M contextual-plus tweets, 85,308 were original tweets with imagery posted by an authoritative source account. Coding was done by the first author and five trained paid coding assistants. The task was binary: an image tweet was coded as either containing a hurricane risk image or not. The coders had a detailed training document with the coding scheme which included 14 types of risk information to be coded as well as example non-hurricane risk images. These categories were developed iteratively across the initial coding task, and then reapplied across the set once the list was stable.

All coding by the assistants was in the presence of the first author. The inter-rater reliability among all six coders was measured with Krippendorff's alpha at three points throughout the five weeks of coding and increased from beginning (α=0.84), to middle (α=0.88), to end (α=0.95); all are considered acceptable for agreement [41]. From this round of coding, we identified 16,789 (19.7% of the original 10M) tweets containing hurricane risk imagery from 500 (62.8% of the original 796) authoritative source accounts.

A second round of coding was conducted by the first author to classify the specific type of each of the identified risk image tweets. The coding scheme was expanded from the 14 categories outlined in Round 1 coding to 22 categories which included greater detail. Any number of the 22 low-level categories could be applied to an image tweet. 258 tweets were removed in this round of coding due to the tweet or account being deleted by the service provider or the user.

In this paper, we selectively scope the analysis for **eight exemplary risk image categories** which are prominent visual hurricane representations spanning multiple types of risks. The Round 2 coding results for these categories are in the first three columns of Table 2. A total of **16,531** (19.4% of all tweets from authoritative sources) original tweets from **489** (61.4% of the initial 796) sources were identified as containing risk imagery. By qualitatively coding these tweets in this iterative manner, we were able to answer **RQ1**.

Table 2: Description of categorized data set of authoritative-source hurricane risk image tweets and their diffusion. Because image categories are non-mutually exclusive and any number of categories can be applied to an image tweet, the columns add up to more than 16,531 (the total number of tweets), or more than 100% for percentages.

Risk Image Category *Non-Mutually Exclusive*	Tweets in Category	% of All Tweets	Replied To	Total Replies	Retweeted	Total Retweets	Quoted	Total Quote Tweets	Not Diffused
Cone of Uncertainty	5103	30.9%	40.9% (2086)	7477	82.7% (4218)	112,503	7.6% (389)	1246	14.9% (761)
Spaghetti	478	2.9%	61.5% (294)	1434	92.3% (441)	8365	5.0% (24)	38	6.9% (33)
Uncertainty	6736	40.8%	42.2% (2844)	10,517	83.2% (5605)	151,209	7.5% (505)	1528	14.5% (978)
Tropical Forecast Advisories	2113	12.8%	38.0% (802)	3372	84.7% (1789)	76,721	9.9% (209)	831	13.7% (289)
Radar/Satellite	6229	37.7%	43.5% (2712)	18,553	85.2% (5310)	283,555	6.1% (377)	1014	13.1% (813)
Evacuation	141	0.9%	44.0% (62)	191	87.9% (124)	3350	6.4% (9)	12	10.6% (15)
NWS/NOAA	3767	22.8%	34.5% (1299)	7250	90.5% (3409)	156,943	9.9% (373)	1119	9.0% (339)
Past	131	0.8%	56.5% (74)	2612	92.4% (121)	34,316	14.5% (19)	94	6.1% (8)

Hurricane Risk Images: What Does (Not) Diffuse

Using the full contextual-plus data of the 9.87M image tweets, we next identified all users who engaged with the 16K original risk image tweets, by either replying directly to them (depth) or retweeting/quote tweeting them, thus creating a new thread around the risk image on their own Twitter timelines (breadth). While reply diffusion reveals a variety of conversations among likely strangers on an authoritative-source image, retweet diffusion can reveal more personalized reactions to images when retweeted on a user's own timeline.

In total, 259,649 unique users (not already included in the 489 authoritative sources) diffused the 16,531 authoritative-source risk image tweets. Of these tweets, 14,184, or 86%, were retweeted a total of 523,509 times by 240,795 users, and 6561 tweets, or 40%, were replied to 36,154 times by 26,692 users. A much smaller proportion (3000, or 7%) were quote-tweeted ("quoted"), while 13% were not diffused at all, meaning they were not replied to, retweeted, nor quoted.

Different hurricane risk image categories diffused differently, shown in Table 2. The most common category is *uncertainty* at 41% of all tweets in the data set; this category comprises mainly *cone of uncertainty* images, which account for 31% of tweets. *Radar/satellite* image tweets, which represent the **current state** (observations) of a storm, were the next most common at 38%. Tweets with a *spaghetti plot*, another representation of the hurricane itself, were uncommon at only 3% of the data; however, this category had the highest percentage of tweets that were replied to at 62% and nearly the highest that were retweeted at 92%, compared to 44% replied to and 85% retweeted for *radar/satellite* and 41% replied to and 83% retweeted for *cone of uncertainty*. Note that cone and spaghetti plot graphics both represent potential **future states** (forecasts) of a storm.

The **past** category consists of imagery conveying information about risks from past hurricanes, whether from days or years ago. Though there were only 131 *past* risk image tweets (0.8% of the data set), these had relatively high rates of replies (56%), retweets (92%), and quote tweets (15%) compared to other categories. This suggests that these images were engaging, and were useful, in some way, to convey risk about the current hurricane threats. We return to this phenomenon of *past* risk imagery in the subsequent studies.

5 STUDY 2: DEFINING & MEASURING DIFFUSION

This study addresses *how* different risk image categories of tweets diffused (**RQ2**). We examine the diffusion patterns of a given tweet or class of tweets relative to image content (the type of risk imagery) and the authoritative sources.

In hurricane events, which require frequent updates as storms evolve, temporal diffusion reveals insights about uptake in relation to physical events. The research therefore considers the period for which risk information stays relevant by investigating how quickly it diffuses, for how long, how many users are reached, and how these diffusion metrics differ across different kinds and sources of risk information.

Method

This study focuses on retweet and reply diffusions. Retweets are often used as a way of constructing networks to understand diffusion [11, 17, 22, 31], and they capture the *breadth* of diffusion of a tweet, in that many users who may or may not follow the original twitterer repost to their own timelines, potentially exposing it to a new set of people. Replies, on the other hand, represent the *depth* of diffusion of a tweet, because they form a discussion thread or threads attached to the original tweet itself.

To measure the diffusion of hurricane risk image tweets based on their retweets and replies, we graphed diffusion "cumulative adoption curves" [17] for individual tweets in the data set, which show cumulative count of retweets or replies across time. These can be thought of as "diffusion signatures" for each tweet. From these signatures, we identified three primary characteristics which contribute to diffusion: total count, duration, and rate. (See examples in Figure 3.) Because of the long tail of tweet diffusion (77.8% of tweets are diffused once or not at all), diffusion analyses using the **count** turned out to be insignificant. Thus we focus on the other two characteristics.

The **duration** represents how long a tweet was diffused, and is measured from the time a tweet was first posted to when it received its final retweet or reply; it represents the "life span" of the tweet. The **rate** is calculated as the gradient of the count of tweets over time, starting from the original tweet. The duration and rate metrics are calculated for retweet-, quote tweet-, and reply-based diffusion. Rate is measured using the first 95% of a tweet's diffusion (i.e. its replies, retweets, or quote tweets), an empirically-determined value which accounts for the steady, initial wave of diffusion and ignores the tail-end where people reply or retweet long after the tweet was first posted, as depicted in Figure 3 in black. The first 95% of the diffusion signature tends to have the steepest slope or, more precisely, the fastest rate of diffusion.

Diffusion by Risk Image Category

For a given risk image category, we conducted statistical analyses (Kruskal-Wallis tests[4]) to compare diffusion of tweets coded with that category to diffusion of all others (i.e., tweets not coded with that category). These tests were performed for each diffusion mechanism (reply, retweet, and quote tweet) and metric (duration and rate). Though we calculated these values for quote tweet diffusion, few differences in quote tweet diffusion metrics were statistically significant, given

[4] For the Kruskal-Wallis test, we compared medians of diffusion distributions.

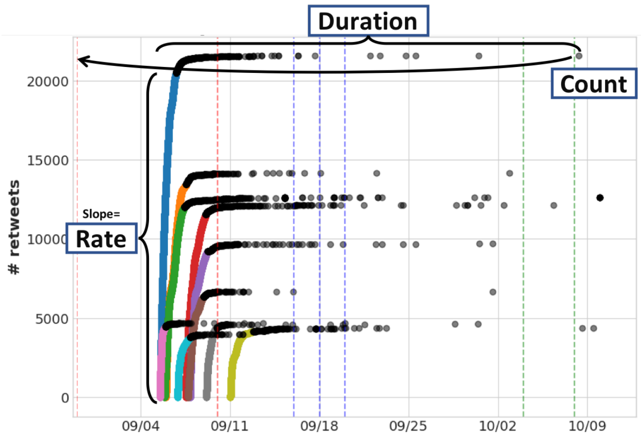

Figure 3: Diffusion adoption curves for top retweeted tweets, as an example to demonstrate metrics. Black portions represent the final 5% of diffusion (e.g., the last 5% of retweets). Rate is calculated for the first 95% of diffusion per tweet; duration and count are calculated for the entire curve.

how infrequently the original 16K tweets were quoted. Additionally, all quote tweet findings that *were* significant were directionally the same as for retweets, i.e., for both retweet and quote tweet diffusion, tweets for a given image category had a longer (shorter) duration, and/or a faster (slower) rate. Thus, only significant results for reply and retweet diffusion are presented in Figure 4, for eight risk image categories.

The **Duration** column of Figure 4 shows for each risk image category the median duration of tweets in that category compared to that of tweets in all other categories, separately for replies and retweets. All significant results reveal longer diffusion durations for tweets with a given risk image category than without.[5] For example, *cone of uncertainty* image tweets are replied to for a significantly longer time (1:14 vs. 0:51) and retweeted for a significantly longer time (2:44 vs. 2:14) than other risk image tweets, as seen by comparing the black bars to the gray bars in the *cone of uncertainty* graph for **Duration**. In fact, all four *forecast* image categories exhibit similar duration patterns of ~1–3 hours for replies and retweets; this mirrors the temporal frequency with which new hurricane forecast information is provided, including by NHC, which updates the cone and other products every three hours when a coastal watch/warning is in effect. The category with the longest duration signature is *past* with a median retweet duration of greater than nine hours, suggesting that the information in these risk images has longer temporal relevance; we explore this further in Study 3.

The **Rate** column of Figure 4 shows for each risk image category the median rate of replies to and retweets of tweets in that category compared to that in all other categories.

[5]This is possible because we only show 8 of the total 22 risk image categories; other categories exhibit the opposite trend, balancing these results.

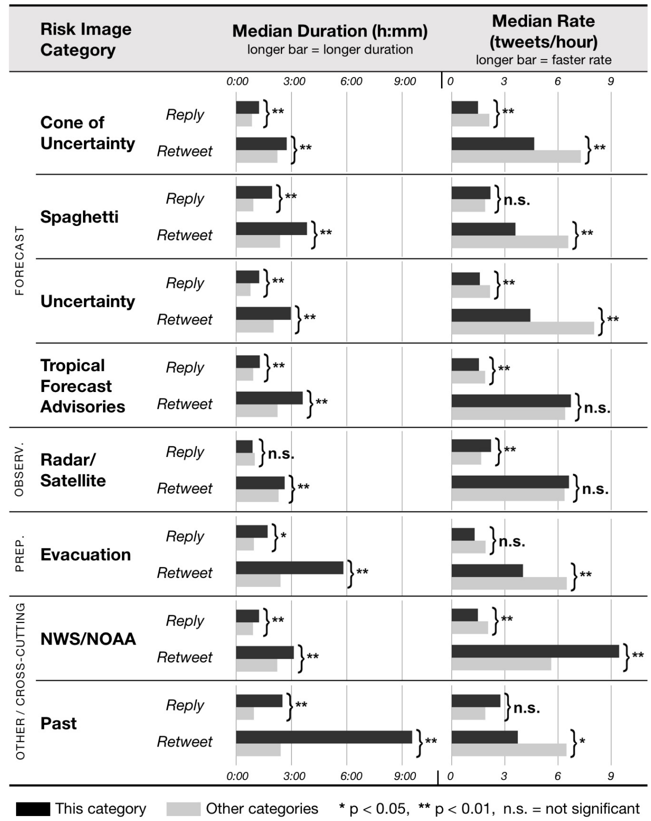

Figure 4: Diffusion measurements (*duration*, the lifespan of a tweet, and *rate*, the speed at which a tweet diffuses) for replies to and retweets of risk image categories.

Radar/satellite images had faster reply diffusion than others (2.2 vs. 1.6 tweets per hour). As a type of *observational* risk image, this category portrays current conditions and very near-term risk in a way that is different from other risk image categories. *Radar/satellite* imagery plausibly has a high rate of diffusion because it conveys what is currently happening in a visually compelling way (e.g. Figure 1c) that people can make meaning out of even if they do not understand the technical details of the image. *NWS/NOAA* tweets—authoritative risk image products stamped with a logo from one of these organizations, regardless of the source user—are retweeted faster (9.4 vs. 5.6 tweets per hour) than other categories and have the fastest retweet rate among the eight presented here. The stamp of a government weather organization on a hurricane risk image likely garners trust from people, prompting them to disseminate these images more quickly than other risk images, at a median rate of one retweet per six minutes.

These findings motivated further qualitative content analysis on how people textually and conversationally respond to different categories to make sense of these sometimes surprising results. This analysis is presented in Section 6.

Diffusion by Authoritative Source Category

As with the risk image categories, we conducted Kruskal-Wallis tests to compare diffusion across source user categories. We additionally implemented a non-parametric pairwise multiple comparisons procedure (Dunn's test) to identify significant differences *among* source user categories, rather than differences based on the presence or absence of a source category. All results have $p < .01$; durations are reported as hh:mm.

Diffusion duration. Risk image tweets from Government accounts have a *longer reply duration* (01:56) than News (00:44) and Weather-News (00:58) tweets. Government tweets and Weather-Other tweets have *longer retweet durations* (05:15 and 04:27, respectively) than all other source categories (News: 02:40, Weather-Government: 02:25, Weather-News: 02:08). This suggests that risk information from government sources might have more long-term relevance for affected people, or more long-term interest for people in general, than risk information shared by news sources.

Diffusion rate. Weather-Other risk image tweets (from independent meteorologists, researchers, and stormchasers) have a *faster reply rate* (3.7 replies per hour) than Weather-News (1.9), Weather-Government (1.3), and Government (1.3) tweets. News tweets have a *faster reply rate* (2.4) than Weather-Government tweets. Weather-Government tweets have the *fastest retweet rate*, (16.2 retweets per hour), while Government tweets have the *slowest retweet rate* (4.1). Though hurricane risk image tweets from government sources diffuse for a longer time than tweets from other sources, they also diffuse more slowly, with news sources and independent weather experts receiving the fastest reply rates, suggesting that the independents have a broader reach when diffusing similar types of information. They may be valuable mediators between the product creators and the general public. Interestingly, though general government tweets are retweeted the slowest, tweets from weather-specific government organizations like the NWS and NHC are retweeted the fastest.

6 STUDY 3: QUALITATIVE ANALYSIS

To understand the meaning behind these differences in diffusion patterns for different risk visuals (e.g. *why* are tweets with the cone of uncertainty replied to for a longer amount of time than other risk image tweets?), and to address **RQ3**, we conducted a targeted qualitative analysis of the textual and conversational aspects of replies and quote tweets.

Method

We analyzed the content of replies to and quote tweets of a sample of the original 16K forecast and risk image tweets. Retweets are not a part of this analysis, as they are simply references to other tweets that do not contain new textual content. We selected one image category per type of visual risk image to get a full representation of the dataset while also scoping the analysis: *cone of uncertainty* for hurricane forecast, *radar/satellite* for observational information, and *evacuation* for preparedness information, as well as *past* as a category that cuts across all the others and was shown in Study 2 to have the longest median duration time. For each category, we sampled up to 50 original tweets that had more than three replies or quote tweets (for some categories, there are fewer than 50 original tweets meeting this criterion). For instance, if tweets in a given category have a longer reply duration than tweets in other categories, then we selected the top 50 tweets with more than three replies in that category with the longest reply durations, then qualitatively analyzed the content of all replies to these 50 tweets. We extracted common high-level themes in relation to diffusion topics, and also remarks that were perhaps not common but noteworthy, in that they suggest alternative insights that englighten this research interpretation, and are noted as such.

Here we report findings from analysis of the textual features of 9291 replies and quote tweets that contributed to diffusion of 253 original, authoritative source risk tweets.

Forecast Risks: Cone of Uncertainty

Cone of uncertainty tweets have longer reply and quote tweet durations and slower reply rates than other categories of hurricane risk images. There is a stark difference in the content of the longest-duration replies compared to that of the longest-duration quote tweets. The replies are primarily grounded in managing uncertainty about risk, i.e. questions about the track *("The million dollar question, where does it turn?")*, requests for more localized information *("@hurricanetrack chances this moves up to US East Coast?")*, and even expert opinions on travel-related risks *("Hey NHC if you had a trip schedule for the BVIs from 9/2-9/8 would you cancel?")*. Quote tweets, on the other hand, are more informal and reactive, especially regarding the onset of Nate after several other hurricanes: *"Break out your bumbershoots and wellies next week!"*, *"Well that's just flippin' fantastic. Ugh."*, *"LEAVE US ALONE! This has been the worst hurricane season!"*.

These differences reflect the mechanics of each type of diffusion: **replies** are attached to the original risk image tweet itself, and thus are potentially exposed to others similarly affected by the hurricane and/or others paying attention to the same tweet, while **quote tweets** are posted only to the quoting user's timeline, and thus are exposed only to the quoting user's friends, family, and followers.

The cone tweets with the slowest reply rates include replies days later that revisit the tweeted forecast information, particularly for Irma (*"So weird to look back at this projection."*), comparing the Irma track to 2016 Hurricane Matthew's (*"This is hurricane Matthew all over again. Everyone is hysterical but nothing will happen yet again"*), and expressing discontent and distrust in the information conveyed in the cone forecast (*"This cone is deceiving. The hurricane force winds are only 50 miles across. This looks like it is 360 miles large"*).

One person replied to a cone tweet 22 hours later saying: *"TRACK HAS GREATLY CHANGED AND IS STILL CHANGING. TAKE THIS TWEET DOWN. It is misleading."* There are several interesting points to note with this reply tweet. First, the speaker is implying that they believe content should be updated in real time even after a postdate, and that any outdated information or information that has since evolved should be taken down, which indicates their misunderstanding about the features of microblog platforms. Second, the speaker is continuing to propagate the tweet by replying to it 22 hours later, meaning others would have seen it appear in their feeds again despite the fact that it was no longer up-to-date. Third, the tweet represents a misunderstanding of the track portrayed by the cone graphic, and that though it is *always* "changing" at any given moment, the static graphics are only released by the NHC at regular intervals (every three hours) and are not updated in real time.

Observational Risks: Radar/Satellite

Radar and satellite image tweets had a faster reply rate than other categories. As mentioned previously, the visually compelling nature of these images likely influences their quick diffusion, and this is supported by many replies such as *"She's beautiful. In a graphic way."*, *"Lovely/terrifying"*, *"Magnificent."* These kinds of comments are not found on any other kind of hurricane risk graphics, thus emphasizing the unique ability of this category of risk imagery to captivate audiences.

People responded to more than the appearance of the graphics, as well. Similar to the cone, people wanted more contextualized information to aid their evacuation decisions (*"Thoughts on if you live in Bradenton??'*) and about the threat to specific areas (*"Where is Irma at right now? I have friends in the USVI."*). Also common were reflections on past hurricanes resembling the current ones (*"Looks like hurricane andrew in 1992, check it out"*) and questions about how to interpret the imagery (*"Maybe you could explain to us laymen what the various colors mean so we can understand better?"*, *"Baseline? What does a normal hurricane look like? Asking for those of us with no expertise on the subject."*).

Preparedness Information: Evacuation

Risk can also be communicated through preparedness and response recommendations, such as evacuation information. Tweets portraying evacuation graphics (typically maps) were replied to for a longer time than other risk image tweets. The replies to such evacuation tweets reveal a strong need for additional and more specific information regarding evacuation orders: people wanted to know whether they should evacuate ("*If I live just north of 6, should I leave?*"), whether evacuations were mandatory ("*Wait, is all of Zone A now mandatory evacuation?*"), where to evacuate to ("*Where are we suppose to go if we are in the evacuation area?!*"), and how to evacuate ("*What is the route that is even open out of new territ. Tell residents where to go*"). Additionally, users expressed that the evacuation maps in tweets lacked timestamps to indicate how up-to-date the information was.

Past Hurricane Graphics

Tweets regarding past hurricanes, whether from years ago or days ago, were both replied to and quote tweeted significantly longer than tweets in other categories. We found this particularly interesting because these tweets do not directly portray current threats in the way other hurricane risk images do. Some of the longest-duration replies and quote tweets compare the past forecast pictured to the current hurricanes at the time to emphasize the latter's threat: *"Andrew was always a small storm. A powerful little buzz saw. Irma, she's husky."* However, the vast majority of diffusion of past forecast tweets is political in nature. Many of these original tweets make reference to climate change, which because of its politicized nature seems to invite this type of response. However, many others were only politicized *by* people's responses in the replies and quote tweets.

The past forecast tweet with both the longest reply and quote tweet durations (12+ days and 2+ days, respectively) compares radar imagery of three hurricanes from 2010 to three in 2017, showing the hurricanes nearly matching in size and location, with the text: *"Absolutely uncanny copy-paste from 7 years ago. Very bizarre. #Irma #Jose #Katia #Igor #Julia #Karl."* The replies are mainly around climate change controversies (*"Not bizarre. They script this shit."*, *"But but but global warming tho"*). Other tweets, such as one from @CNN stating *"Yes, climate change made Harvey and Irma worse,"* clearly invite such politicized responses. The surprising diffusion of past hurricane image tweets is thus attributed more to people's desire to discuss politics than to anything having to do with the hurricanes or their risks per se.

Political and Off-Topic Diffusion

These qualitative analyses uncovered many replies and quote tweets that were not related to hurricanes, but rather about fake climate change, there not being an appointed FEMA administrator at the time, and more, and were intended to make political statements. The tweet with the most replies (2825) and third most retweets (13,988) was an urgent, yet non-controversial hurricane risk message regarding Hurricane Harvey: *"NOTICE: The levee at Columbia Lakes has been breached!! GET OUT NOW!!"* The reason for its vast diffusion was that it had been retweeted by @realDonaldTrump, the official Twitter handle for the current US President, thus exposing it to his 3.7M followers. More than 75% of replies to this tweet were unrelated to hurricanes, and were instead commentary on Trump and US politics. This was the only one of the 16,531 tweets in the dataset diffused by Trump.

To determine whether these political tweets were a large percentage of replies overall, we randomly sampled and read $n = 100$ tweets from the 36K total replies to risk images to classify each as purely political (off-topic) or having anything to do with hurricanes or forecasting. We found 25 off-topic tweets. Keywords derived from the data helped in identifying other off-topic tweets: "trump", "climate", "climate change", "global", "warming", "heating", "obama", "antichrist", "fema", "donald", "drumpf", "white", "suprem", "fake", "bush." These terms were used only as a starting point, as not all off-topic tweets contain one, and not all tweets containing one were off-topic. Based on this, a sample of $n = 500$ ensured no more than a ±2.5% error in our off-topic prevalence estimate for the entire dataset; in this sample, we found 105 off-topic tweets for an estimated prevalence of 21%. Excluding from this sample the off-topic replies to the tweet above that was retweeted by Trump, this decreases to 17.3%.

7 DISCUSSION AND IMPLICATIONS

This multi-method investigation uses social media as the platform for examination of 1) what risk images are shared for hurricanes, 2) how these images differently diffuse based on type of risk portrayed and type of authoritative source user, and 3) why these images diffuse the way they do based on recipients' responses and questions. In this section, we discuss how the findings contribute to HCI, highlighting implications for design, policy, and methodology.

Design Implications

We might think of current risk imagery as boundary objects that sit between the scientific, practitioner, and lay communities—and therefore bear a great deal of burden. Because risk communication must be directed to often millions of people under threat, the implications of risk interpretation are many: they must be put in relation to both the assessment of danger felt by any one person as well as the various costs associated with mass response. The scientific representations present authoritative assessment, but cannot on their own resolve what people consequently do. This depends on further translational work downstream, which is where there are opportunities for the innovation of new information products to support risk communication.

Thus, an important implication from this work is how to reduce the burden on risk imagery for communicating critical, yet often uncertain, information to various stakeholders. Study 3 showed that with all hurricane risk visualizations, laypeople/the public request information that is localized to their particular situation, whether it pertains to where they live, where they might travel, or where their friends or family are located. **Risk information visualizations should be designed such that people may contextualize risks to their own situations and utilize the information more effectively**, rather than interpret risks only as they apply broadly to an entire state or geographic region. Media platforms could enable the use of interactive images that allow users to zoom and pan to more granular levels. Authoritative sources who generate and share these images could additionally outline forecasted impacts at the city- or neighborhood-level to be visually incorporated into the risk image or included in the associated text post. Further still, distilled, localized information for various locations could be presented in *new* risk representations meant specifically for public use, rather than combining public risk information with complex, scientific representations.

Such visualizations of uncertainty, risk, and probability that are consumed by the public are notoriously difficult to both render and interpret [19, 25]. However, uncertainty in visualizations can help people to make better estimates [18], suggesting that uncertainty should be made *more* obvious in hurricane risk images. The NHC track forecast cone is updated annually, but graphics like spaghetti plots are not standardized and often elicit confusion from the public. **Risk representations should convey uncertainty as appropriate in understandable, meaningful ways so that people can make best use of the information in interpreting risk.**

Methodological Contributions

This research demonstrates the complexities of studying the diffusion of large-scale phenomena—in this case hurricane risk—on widely adopted social network platforms. As hurricanes rapidly evolve and information about them is regularly updated, people must regularly (re)orient to the changing risk. This rapidly changing information delivery and consumption environment that we see in this case, but exists in others too, thus requires specialized methods to describe diffusion.

As an HCI contribution, we empirically determined in Study 2 that the diffusion of tweets could be quantified by two characteristics—duration and rate—in terms of three different mechanisms—retweets, quote tweets, and replies—to identify those image tweets that possess distinctive diffusion signatures. This quantitative analysis in turn inspired qualitative analysis to obtain further insight about why certain images and image categories diffuse. By examining replies back to authoritative sources as well as quote tweets to friends and followers, we gain insight on how people interpret risk imagery through their questions and comments, as well as the different uses of replies and quote tweets for interacting with risk images. We see reactions that are contextualized relative to people's own hazards risk situations as well as to their worldviews about other matters that disasters encompass, such as reactions to climate change arguments and political disputes around emergency response activities. **This research considers risk image diffusion as more than just risk assessment in a strictly rationalized, information retrieval sense, and thus assesses reasons why certain threads endure as part of a complex information environment**, building upon and contributing to the field of crisis informatics research by deepening the questions we pose and the approaches we take.

Policy Implications

A final aim of this research is to support the needs of the weather scientific and practitioner communities to understand socio-behavioral phenomena around their information products. To highlight a few points: As evolving atmospheric phenomena, hurricanes can strengthen, weaken, shift course, slow down or stall. Accordingly, the risks posed by them evolve as well, and this evolution is represented in the risk information produced. Updated observations and forecasts of the hurricanes are then provided on a regular basis, and evacuation orders are issued or rescinded. The diffusion signature patterns reveal a correspondence between the duration and the temporal "relevance" of the information produced about evolving risks. This suggests that **people are broadly attuned to the temporal "cadence" of different types of hurricane risk information**—i.e., that forecasts and observations turn over on the order of 1 to 3 hours, evacuation information is more static, and that past hurricane information can be indefinitely "relevant" as an indicator of potential risks of a current hurricane. The inclusion of timestamps on all risk imagery would further support this.

An important exception arose for the critical category of evacuation, which provides actionable risk information for the public. The median duration for retweets of evacuation orders is about 6 hours, even though orders are issued further in advance of landfall and not usually updated. This suggests that **evacuation information should be regularly reintroduced into the social media sphere for it to receive timely attention.**

Additionally, this research suggests potential changes to practices around sharing hurricane risk imagery on social media, particularly by authoritative sources, supporting related HCI work on rioting behavior [14] and emergency communication from police and fire [24]. Studies 1 and 2 identified and quantified diffusion of several types of hurricane risk images, but Study 3 showed that such diffusion metrics may not tell the full story, as diffusion is an indicator of many different reactions and responses. Thus, in addition to using quantifiable diffusion statistics offered by social media platforms, such as "reach" and "impressions" from Twitter and Facebook, **authoritative sources who produce and share risk information should be heavily engaged in the resulting conversations** to answer questions, clear confusion, and gauge and shape the public's understanding of risk.

8 ACKNOWLEDGMENTS

We thank our colleagues at the University of Colorado Boulder and the National Center for Atmospheric Research for their assistance on this paper and related research that informed this work. These include Jennings Anderson, Kenneth Mark Anderson, Kevin Ash, Gerard Casas Saez, Jen Henderson, Heather Lazrus, Rebecca Morss, and Olga Wilhelmi. We are especially grateful to our undergraduate research assistants: Jamie Ackerson, Hande Batan, Jason Cummings, Marissa Kelley, Rohan Ramnani, and Steven Yatko. This research was funded by NSF grants 1564275 and 1331490, and an NSF GRFP.

REFERENCES

[1] Firoj Alam, Ferda Ofli, Muhammad Imran, and Michael Aupetit. 2018. A Twitter Tale of Three Hurricanes: Harvey, Irma, and Maria. In *Proceedings of the 15th International Conference on Information Systems for Crisis Response and Management (ISCRAM)*. https://www.unocha.org/legacy/what-we-do/coordination-tools/cluster-coordinationhttp://arxiv.org/abs/1805.05144

[2] Jennings Anderson, Gerard Casas Saez, Kenneth M. Anderson, Leysia Palen, and Rebecca Morss. in press. Incorporating Context and Location Into Social Media Analysis: A Scalable, Cloud-Based Approach for More Powerful Data Science. In *Proceedings of the 52nd Annual Hawaii International Conference on Systems Sciences*.

[3] Jennings Anderson, Marina Kogan, Melissa Bica, Leysia Palen, Kenneth Anderson, Kevin Stowe, Rebecca Morss, Julie Demuth, Heather Lazrus, Olga Wilhelmi, and Jennifer Henderson. 2016. Far Far Away in Far Rockaway: Responses to Risks and Impacts during Hurricane Sandy through First-Person Social Media Narratives. In *Proceedings of the 13th International Conference on Information Systems for Crisis Response and Management (ISCRAM âĂŸ16)*.

[4] Cynthia A Andrews, Elodie S Fichet, Yuwei Ding, Emma S Spiro, and Kate Starbird. 2016. Keeping Up with the Tweet-dashians: The Impact of 'Official' Accounts on Online Rumoring. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing - CSCW '16*. ACM Press, New York, New York, USA, 451–464. https://doi.org/10.1145/2818048.2819986

[5] Eytan Bakshy, Itamar Rosenn, Cameron Marlow, and Lada Adamic. 2012. The Role of Social Networks in Information Diffusion. In *Proceedings of the 21st international conference on World Wide Web*. 519–528. https://doi.org/10.1145/2187836.2187907

[6] Melissa Bica, Leysia Palen, and Chris Bopp. 2017. Visual Representations of Disaster. In *Proceedings of the ACM 2017 Conference on Computer Supported Cooperative Work*. ACM Press, New York, New York, USA, 1262–1276. https://doi.org/10.1145/2998181.2998212

[7] Ann Bostrom, Rebecca E. Morss, Jeffrey K. Lazo, Julie L. Demuth, Heather Lazrus, and Rebecca Hudson. 2016. A Mental Models Study of Hurricane Forecast and Warning Production, Communication, and Decision-Making. *Weather, Climate, and Society* 8, 2 (4 2016), 111–129. https://doi.org/10.1175/WCAS-D-15-0033.1

[8] Kenneth Broad, Anthony Leiserowitz, Jessica Weinkle, and Marissa Steketee. 2007. Misinterpretations of the "cone of uncertainty" in Florida during the 2004 hurricane season. *Bulletin of the American Meteorological Society* 88, 5 (5 2007), 651–667. https://doi.org/10.1175/BAMS-88-5-651

[9] Axel Bruns and Jean Burgess. 2012. Local and global responses to disaster: #eqnz and the Christchurch earthquake. In *Disaster and Emergency Management Conference*, Peter Sugg (Ed.). 86–103.

[10] Justin Cheng, Lada Adamic, P Alex Dow, Jon Michael Kleinberg, and Jure Leskovec. 2014. Can cascades be predicted?. In *Proceedings of the 23rd international conference on World wide web*. 925–936. https://doi.org/10.1145/2566486.2567997

[11] Munmun De Choudhury. 2015. *YouâĂŹre Happy, IâĂŹm Happy: Diffusion of Mood Expression on Twitter*. Hanbit Media, Inc. 169–179 pages. https://dl.acm.org/citation.cfm?id=2729546

[12] Julie L. Demuth, Rebecca E. Morss, Jeffrey K. Lazo, and Craig Trumbo. 2016. The Effects of Past Hurricane Experiences on Evacuation Intentions through Risk Perception and Efficacy Beliefs: A Mediation Analysis. *Weather, Climate, and Society* 8, 4 (10 2016), 327–344. https://doi.org/10.1175/WCAS-D-15-0074.1

[13] Julie L. Demuth, Rebecca E. Morss, Betty Hearn Morrow, and Jeffrey K. Lazo. 2012. Creation and communication of hurricane risk information. *Bulletin of the American Meteorological Society* 93, 8 (8 2012), 1133–1145. https://doi.org/10.1175/BAMS-D-11-00150.1

[14] Sebastian Denef, Petra S Bayerl, and Nico Kaptein. 2013. *Social Media and the Police-Tweeting Practices of British Police Forces during the August 2011 Riots*. https://www.fit.fraunhofer.de/content/dam/fit/de/documents/ukriots%20v89.pdf

[15] Gina Eosco. 2008. *A Study of Visual Communication: Cyclones, Cones, and Confusion*. Ph.D. Dissertation.

[16] Adrien Friggeri, La Adamic, Dean Eckles, and Justin Cheng. 2014. Rumor Cascades. In *Proceedings of the Eighth International AAAI Conference on Weblogs and Social Media*. 101–110. http://www.aaai.org/ocs/index.php/ICWSM/ICWSM14/paper/download/8122/8110

[17] Sharad Goel, Ashton Anderson, Jake Hofman, and Duncan J. Watts. 2016. The Structural Virality of Online Diffusion. *Management Science* 62, 1 (7 2016), 180–196. https://doi.org/10.1287/mnsc.2015.2158

[18] Miriam Greis, Aditi Joshi, Ken Singer, Albrecht Schmidt, Tonja Machulla, and L M U Munich. 2018. Uncertainty Visualization Influences how Humans Aggregate Discrepant Information. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems - CHI '18*. ACM Press, New York, New York, USA, 1–12. https://doi.org/10.1145/3173574.3174079

[19] Miriam Greis, Thorsten Ohler, Niels Henze, and Albrecht Schmidt. 2015. Investigating Representation Alternatives for Communicating Uncertainty to Non-experts. In *Human-Computer Interaction – INTERACT 2015*, Julio Abascal, Simone Barbosa, Mirko Fetter, Tom Gross, Philippe Palanque, and Marco Winckler (Eds.). Springer International Publishing, 256–263. https://doi.org/10.1007/978-3-319-22723-8{_}21

[20] Xinning Gui, Yubo Kou, Kathleen Pine, Elisa Ladaw, Harold Kim, Eli Suzuki-Gill, and Yunan Chen. 2018. Multidimensional Risk Communication: Public Discourse on Risks during an Emerging Epidemic. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM Press, New York, New York, USA, 214:1–214:14. https://doi.org/10.1145/3173574.3173788

[21] Xinning Gui, Yubo Kou, Kathleen H. Pine, and Yunan Chen. 2017. Managing Uncertainty: Using Social Media for Risk Assessment during a Public Health Crisis. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems - CHI '17*. ACM Press, New York, New York, USA, 4520–4533. https://doi.org/10.1145/3025453.3025891

[22] Scott A. Hale and Scott A. 2014. Global connectivity and multilinguals in the Twitter network. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems - CHI '14*. ACM Press, New York, New York, USA, 833–842. https://doi.org/10.1145/2556288.2557203

[23] Samiul Hasan, Satish Ukkusuri, Hugh Gladwin, and Pamela Murray-Tuite. 2011. Behavioral Model to Understand Household-Level Hurricane Evacuation Decision Making. *Journal of Transportation Engineering* 137, 5 (2011), 341–348. https://doi.org/10.1061/(ASCE)TE.1943-5436.0000223

[24] Amanda L. Hughes, Lise A. A. St. Denis, Leysia Palen, Kenneth M. Anderson, Amanda L. Hughes, Lise A. A. St. Denis, Leysia Palen, and Kenneth M. Anderson. 2014. Online public communications by police & fire services during the 2012 Hurricane Sandy. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems - CHI '14*. ACM Press, New York, New York, USA, 1505–1514. https://doi.org/10.1145/2556288.2557227

[25] Harald Ibrekk and M. Granger Morgan. 1987. Graphical Communication of Uncertain Quantities to Nontechnical People. *Risk Analysis* 7, 4 (12 1987), 519–529. https://doi.org/10.1111/j.1539-6924.1987.tb00488.x

[26] Muhammad Imran, Carlos Castillo, Fernando Diaz, and Sarah Vieweg. 2014. Processing Social Media Messages in Mass Emergency: A Survey. *ACM Comput. Surv.* 47, 4 (2014), 67:1–67:38. https://doi.org/10.1145/2771588

[27] Muhammad Imran, Prasenjit Mitra, and Carlos Castillo. 2016. Twitter as a Lifeline: Human-annotated Twitter Corpora for NLP of Crisis-related Messages. In *Proceedings of the 10th Language Resources and Evaluation Conference (LREC)*. 1638–1643. https://en.wikipedia.org/wiki/Receiver_http://arxiv.org/abs/1605.05894

[28] Susan Joslyn and Sonia Savelli. 2010. Communicating forecast uncertainty: Public perception of weather forecast uncertainty. *Meteorological Applications* 17, 2 (4 2010), 180–195. https://doi.org/10.1002/met.190

[29] Brian Keegan, Darren Gergle, and Noshir Contractor. 2011. Hot off the wiki: dynamics, practices, and structures in Wikipedia's coverage of the Tōhoku catastrophes. In *Proceedings of the 7th International Symposium on Wikis and Open Collaboration*. 105–113. https://doi.org/10.1145/2038558.2038577

[30] Marina Kogan, Leysia Palen, and Kenneth M. Anderson. 2015. Think Local, Retweet Global. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing - CSCW '15*. ACM Press, Vancouver, 981–993. https://doi.org/10.1145/2675133.2675218

[31] Kristina Lerman and Rumi Ghosh. 2010. Information Contagion: an Empirical Study of the Spread of News on Digg and Twitter Social Networks. In *Association for the Advancement of Artificial Intelligence*. 90–97. https://doi.org/10.1146/annurev.an.03.100174.001431

[32] Michael K. Lindell and Ronald W. Perry. 2012. The Protective Action Decision Model: Theoretical Modifications and Additional Evidence. *Risk Analysis* 32, 4 (4 2012), 616–632. https://doi.org/10.1111/j.1539-6924.2011.01647.x

[33] Isaac M. Lipkus and J G Hollands. 1999. The visual communication of risk. *Journal of the National Cancer Institute. Monographs* 27701, 25 (1 1999), 149–63. https://doi.org/10.1093/oxfordjournals.jncimonographs.a024191

[34] Alan M. MacEachren, Anthony Robinson, Susan Hopper, Steven Gardner, Robert Murray, Mark Gahegan, and Elisabeth Hetzler. 2005. Visualizing Geospatial Information Uncertainty: What We Know and What We Need to Know. *Cartography and Geographic Information Science* 32, 3 (2005), 139–160. http://test.scripts.psu.edu/students/a/c/acr181/MacEachren_et_al_FINAL.pdf

[35] Jim Maddock, Kate Starbird, Haneen J. Al-Hassani, Daniel E. Sandoval, Mania Orand, and Robert M. Mason. 2015. Characterizing Online Rumoring Behavior Using Multi-Dimensional Signatures. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing - CSCW '15*. ACM Press, New York, New York, USA, 228–241. https://doi.org/10.1145/2675133.2675280

[36] Marcelo Mendoza, Barbara Poblete, and Carlos Castillo. 2010. Twitter Under Crisis : Can we trust what we RT?. In *1st Workshop on Social Media Analytics (SOMA '10)*, Vol. 1060. ACM Press, New York, New York, USA, 71–79. https://doi.org/10.1145/1964858.1964869

[37] Rebecca E Morss, Julie L Demuth, and Jeffrey K Lazo. 2008. Communicating Uncertainty in Weather Forecasts: A Survey of the U.S. Public. *Weather and Forecasting* 23, 5 (2008), 974–991. https://doi.org/10.1175/

2008WAF2007088.1
[38] Rebecca E. Morss, Julie L. Demuth, Heather Lazrus, Leysia Palen, C. Michael Barton, Christopher A. Davis, Chris Snyder, Olga V. Wilhelmi, Kenneth M. Anderson, David A. Ahijevych, Jennings Anderson, Melissa Bica, Kathryn R. Fossell, Jennifer Henderson, Marina Kogan, Kevin Stowe, Joshua Watts, Rebecca E. Morss, Julie L. Demuth, Heather Lazrus, Leysia Palen, C. Michael Barton, Christopher A. Davis, Chris Snyder, Olga V. Wilhelmi, Kenneth M. Anderson, David A. Ahijevych, Jennings Anderson, Melissa Bica, Kathryn R. Fossell, Jennifer Henderson, Marina Kogan, Kevin Stowe, and Joshua Watts. 2017. Hazardous Weather Prediction and Communication in the Modern Information Environment. *Bulletin of the American Meteorological Society* 98, 12 (12 2017), 2653–2674. https://doi.org/10.1175/BAMS-D-16-0058.1
[39] Rebecca E. Morss, Jeffrey K. Lazo, and Julie L. Demuth. 2010. Examining the use of weather forecasts in decision scenarios: Results from a US survey with implications for uncertainty communication. *Meteorological Applications* 17, 2 (5 2010), 149–162. https://doi.org/10.1002/met.196
[40] Tara M. Mortensen, Kevin Hull, and Kelli S. Boling. 2017. Really Social Disaster: An Examination of Photo Sharing on Twitter During the #SCFlood. *Visual Communication Quarterly* 24, 4 (10 2017), 219–229. https://doi.org/10.1080/15551393.2017.1388704
[41] Kimberly A Neuendorf. 2016. *The content analysis guidebook.* Sage.
[42] National Academies of Sciences, Engineering, and Medicine. 2018. *Integrating Social and Behavioral Sciences Within the Weather Enterprise.* The National Academies Press, Washington, D.C. https://doi.org/10.17226/24865
[43] Lace M. Padilla, Ian T. Ruginski, and Sarah H. Creem-Regehr. 2017. Effects of ensemble and summary displays on interpretations of geospatial uncertainty data. *Cognitive Research: Principles and Implications* 2, 1 (12 2017), 40. https://doi.org/10.1186/s41235-017-0076-1
[44] Leysia Palen, Kenneth M. Anderson, Gloria Mark, James Martin, Douglas Sicker, Martha Palmer, and Dirk Grunwald. 2010. A vision for technology-mediated support for public participation & assistance in mass emergencies & disasters. *Proceedings of the 2010 ACMBCS Visions of Computer Science Conference* (4 2010), 1âĂŞ12. http://dl.acm.org/citation.cfm?id=1811182.1811194http://portal.acm.org/citation.cfm?id=1811182.1811194
[45] Florian Pappenberger, Elisabeth Stephens, Jutta Thielen, Peter Salamon, David Demeritt, Schalk Jan van Andel, Fredrik Wetterhall, and Lorenzo Alfieri. 2013. Visualizing probabilistic flood forecast information: Expert preferences and perceptions of best practice in uncertainty communication. *Hydrological Processes* 27, 1 (1 2013), 132–146. https://doi.org/10.1002/hyp.9253
[46] Laura N. Rickard, Jonathon P. Schuldt, Gina M. Eosco, Clifford W. Scherer, and Ricardo A. Daziano. 2017. The Proof is in the Picture: The Influence of Imagery and Experience in Perceptions of Hurricane Messaging. *Weather, Climate, and Society* 9 (2017), 471–485. https://doi.org/10.1175/WCAS-D-16-0048.1
[47] Simon Rogers. 2014. What fuels a Tweet's engagement? (2014). https://blog.twitter.com/2014/what-fuels-a-tweets-engagement
[48] Daniel M. Romero, Brendan Meeder, and Jon Kleinberg. 2011. Differences in the mechanics of information diffusion across topics. In *Proceedings of the 20th international conference on World wide web - WWW '11.* ACM Press, New York, New York, USA, 695. https://doi.org/10.1145/1963405.1963503
[49] Florian Roth. 2012. *Visualizing Risk: The Use of Graphical Elements in Risk Analysis and Communications.* Technical Report July.
[50] Ian T. Ruginski, Alexander P. Boone, Lace M. Padilla, Le Liu, Nahal Heydari, Heidi S. Kramer, Mary Hegarty, William B. Thompson, Donald H. House, and Sarah H. Creem-Regehr. 2016. Non-expert interpretations of hurricane forecast uncertainty visualizations. *Spatial Cognition & Computation* 16, 2 (4 2016), 154–172. https://doi.org/10.1080/13875868.2015.1137577
[51] Hyunjin Seo. 2014. Visual Propaganda in the Age of Social Media: An Empirical Analysis of Twitter Images During the 2012 Israeli-Hamas Conflict. *Visual Communication Quarterly* 21, June (11 2014), 150–161. https://doi.org/10.1080/15551393.2014.955501
[52] Kate Starbird, Dharma Dailey, Owla Mohamed, Gina Lee, and Emma S. Spiro. 2018. Engage Early, Correct More : How Journalists Participate in False Rumors Online during Crisis Events. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems - CHI '18.* ACM Press, New York, New York, USA, 1–12. https://doi.org/10.1145/3173574.3173679
[53] Kate Starbird and Leysia Palen. 2010. Pass It On?: Retweeting in Mass Emergency. *Proceedings of the 7th International ISCRAM Conference* December 2004 (2010), 1–10. https://doi.org/10.1111/j.1556-4029.2009.01231.x
[54] Kate Starbird and Leysia Palen. 2011. "Voluntweeters": Self-Organizing by Digital Volunteers in Times of Crisis. In *Proceedings of the 2011 annual conference on Human factors in computing systems (CHI '11).* ACM Press, New York, New York, USA, 1071–1080. https://doi.org/10.1145/1978942.1979102
[55] Kate Starbird and Leysia Palen. 2012. (How) will the revolution be retweeted?: information diffusion and the 2011 Egyptian uprising. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work - CSCW '12.* 7âĂŞ16. https://doi.org/10.1145/2145204.2145212
[56] Kate Starbird, Leysia Palen, Amanda L. Hughes, and Sarah Vieweg. 2010. Chatter on The Red: What Hazards Threat Reveals about the Social Life of Microblogged Information. In *Proceedings of the 2010 ACM conference on Computer supported cooperative work - CSCW '10.* ACM Press, New York, New York, USA, 241. https://doi.org/10.1145/1718918.1718965
[57] Susanne Tak, Alexander Toet, and Jan van Erp. 2015. Public Understanding of Visual Representations of Uncertainty in Temperature Forecasts. *Journal of Cognitive Engineering and Decision Making* 9, 3 (9 2015), 241–262. https://doi.org/10.1177/1555343415591275
[58] Farida Vis, Simon Faulkner, Katy Parry, Yana Manyukhina, Lisa Evans, Yana Manykhina, and Lisa Evans. 2013. Twitpic-ing the riots: analysing images shared on Twitter during the 2011 UK riots. In *Twitter and Society*, Katrin Weller, Axel Bruns, Jean Burgess, Merja Mahrt, and Cornelius Puschmann (Eds.). Number July 2005. New York:Peter Lang, New York, New York, USA, 385–398. http://eprints.whiterose.ac.uk/79098/8/WRRO_79098.pdf
[59] Katrin Weller, Axel Bruns, Jean Burgess, Merja Mahrt, and Cornelius Puschmann. 2014. *Twitter and Society.* 1–486 pages. https://doi.org/10.1017/CBO9781107415324.004
[60] Joanne I. White, Leysia Palen, and Kenneth M. Anderson. 2014. Digital mobilization in disaster response: The Work & Self-Organization of On-LIne Pet Advocates in Response to Hurricane Sandy. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing - CSCW '14.* ACM Press, New York, New York, USA, 866–876. https://doi.org/10.1145/2531602.2531633