

Characterizing Collective Attention via Descriptor Context: A Case Study of Public Discussions of Crisis Events

Ian Stewart, Diyi Yang, Jacob Eisenstein

Georgia Institute of Technology
School of Interactive Computing
Atlanta, GA 30308

{istewart6, jacob}@gatech.edu, diyi.yang@cc.gatech.edu

Abstract

Social media datasets make it possible to rapidly quantify *collective attention* to emerging topics and breaking news, such as crisis events. Collective attention is typically measured by aggregate counts, such as the number of posts that mention a name or hashtag. But according to rationalist models of natural language communication, the collective salience of each entity will be expressed not only in how often it is mentioned, but in the form that those mentions take. This is because natural language communication is premised on (and customized to) the expectations that speakers and writers have about how their messages will be interpreted by the intended audience. We test this idea by conducting a large-scale analysis of public online discussions of breaking news events on Facebook and Twitter, focusing on five recent crisis events. We examine how people refer to locations, focusing specifically on contextual descriptors, such as “*San Juan*” versus “*San Juan, Puerto Rico*.” Rationalist accounts of natural language communication predict that such descriptors will be unnecessary (and therefore omitted) when the named entity is expected to have high prior salience to the reader. We find that the use of contextual descriptors is indeed associated with proxies for social and informational expectations, including macro-level factors like the location’s global salience and micro-level factors like audience engagement. We also find a consistent decrease in descriptor context use over the lifespan of each crisis event. These findings provide evidence about how social media users communicate with their audiences, and point towards more fine-grained models of collective attention that may help researchers and crisis response organizations to better understand public perception of unfolding crisis events.

1 Introduction

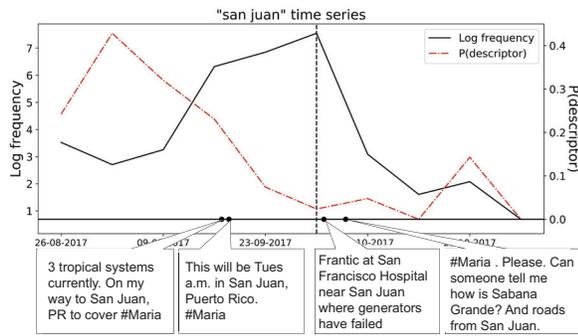
Breaking news events, such as crises, can attract significant *collective attention* from the general public (Lin et al. 2014), resulting in bursts of discussion on social media (Leavitt and Clark 2014; Lehmann et al. 2012). During such events, attention is often focused on entities (e.g., people, locations, and organizations) that are highly relevant to the unfolding event (Wakamiya et al. 2015). A spike in attention directed

toward a particular entity may signal an important update, such as the need for aid for the location (Varga et al. 2013).

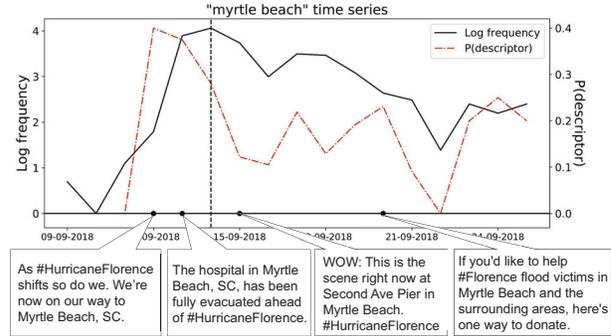
While collective attention is often measured with activity metrics such as post volume (Mitra, Wright, and Gilbert 2016), we propose to characterize collective attention not by *how often* an entity is mentioned, but by *how* it is mentioned. Content-based metrics can be more accurate, as they are less sensitive to data sparsity and biases, such as when a crisis limits internet accessibility. Furthermore, measuring the content of collective attention can provide insight into writer’s expectations of reader knowledge, which may become visible as they adapt their writing to address a local or general audience.

To understand how the content of entity references can change with collective attention, consider Hurricane Maria, which struck Puerto Rico in September 2017. As more people became familiar with the locations mentioned in news coverage about the island (DiJulio, Muñana, and Brodie 2017), news headlines and articles referred to “*San Juan*” without extra contextual descriptors such as “*the capital of Puerto Rico*.” This is consistent with a rational model of communication in which linguistic contextualization is used for entities that might otherwise be unknown or ambiguous (Prince 1992; Staliūnaitė et al. 2018): as San Juan became increasingly salient through repeated mentions, readers could be expected to understand the reference without additional context. Figure 1 presents evidence from Twitter in favor of this hypothesis: the locations of San Juan (1a) and Myrtle Beach (1b) received fewer contextualizing descriptors following peaks in the volume of mentions in discussions of Hurricanes Maria and Florence respectively.

While global salience plays an important role, more fine-grained factors are also at work. Authors may add or remove context information based on their expectations about their specific audience and on the availability of additional context such as hyperlinks to external articles. Furthermore, even if we observe an aggregate change in collective attention, we cannot be sure whether the trend is due to change in the author population or a change in the behavior of individual authors. We disentangle these macro-level and micro-level factors in a multivariate



(a) Timeline for mentions of “San Juan” during Hurricane Maria, with example tweets below.



(b) Timeline for mentions of “Myrtle Beach” during Hurricane Florence, with example tweets below.

Figure 1: Example of collective attention expressed toward location mentions in discussion of various hurricanes on Twitter. Left y-axis (black solid line) indicates the location’s log frequency, right y-axis (red dotted line) indicates the location’s probability of receiving a descriptor phrase such as “San Juan, Puerto Rico.” For example, a 25% probability for “San Juan” means that 25% of all mentions of “San Juan” had a descriptor phrase.

analysis of descriptor phrase usage in discussions of crisis events, using data from public discussions of five recent natural disasters on Facebook and Twitter. We investigate the usage of contextual descriptor phrases in references to locations affected by hurricanes, which we link to various proxies for information expectations: temporal trends in relation to the event itself; properties of the author, audience, and entity; and the presence of extra-linguistic context such as hyperlinks and images. This investigation addresses the following research questions:

- **RQ1:** What factors influence the use of descriptor context in reference to locations of hurricane events?
- **RQ2a:** How does the use of descriptor context for locations change over time at a collective level?
- **RQ2b:** How does the use of descriptor context for locations change at an individual author level?

We now briefly summarize the high-level findings. In a dataset of Facebook posts from public groups concerning Hurricane Maria relief, we find that location mentions received descriptors more often when the locations were not local to the group of discussion, suggesting that descriptors may be used to help explain new information to audiences. In a dataset of public Twitter posts related to five hurricane events, we find that the aggregate rate of descriptor phrases decreased following the peaks in these locations’ collective attention, supporting prior findings on named entity references in professional newstext (Staliūnaitė et al. 2018).

However, this result is supplemented by more fine-grained effects that also support a rational account of entity reference: authors used fewer descriptors if they had mentioned a location before, and more descriptors if their previous posts received high audience engagement (as measured by retweets and likes), which suggest a larger (and potentially less contextualized) audience. In sum, this work identifies strong connections between shared information

and entity references, which supplements existing linguistic theory while offering researchers and practitioners new tools for measuring and understanding collective attention in crisis events. All code is available for use.¹

2 Related Work

The term *collective attention* refers to the degree to which public interest and awareness is focused on individual events, entities, or topics (Sasahara et al. 2013). Collective attention can shift either rapidly or gradually (Wu and Huberman 2007), often in response to large-scale events such as sports games (Lehmann et al. 2012), natural disasters (Varga et al. 2013), and political controversy (Garimella et al. 2017). Prior studies of social media have often quantified collective attention using the volume of posting and sharing activity (Leavitt and Clark 2014; Mitra, Wright, and Gilbert 2016). We supplement this prior work by focusing on how linguistic content changes with collective attention over crisis events.

When referring to an entity such as a location, the writer may add descriptive information in the form of a dependent clause (Kang et al. 2019). The dependent clause may describe attributes of the entity that are relevant to a specific topic, such as “San Juan, epicenter of the Hurricane Maria relief effort,” or attributes that are generally relevant, such as “San Juan, Puerto Rico.” Rationalist models of communication (Grice 1975; Prince 1992) predict that such descriptors will be particularly necessary for entities that are not already salient to the audience. These predictions have received empirical support in corpus analyses: Siddharthan, Nenkova, and McKeown (2011) found that highly *salient* or *important* named entities in a summary of a news story are more likely to be understood as shared knowledge and are therefore less likely to need a descriptor phrase. Similarly,

¹Repository here: https://github.com/ianbstewart/collective_attention.

Event	Hashtags	Date range	Tweets	LOCATION NEs	LOCATION examples
Florence	#florence, #hurricaneflorence	[30-08-18, 26-09-18]	66595	28670	Wilmington, New Bern, Myrtle Beach
Harvey	#harvey, #hurricaneharvey	[17-08-17, 10-09-17]	679400	181636	Houston, Corpus Christi, Rockport
Irma	#irma, #hurricaneirma	[29-08-17, 20-09-17]	809423	229315	Miami, Tampa, Naples
Maria	#maria, #hurricanemaria, #huracanmaria	[15-09-17, 09-10-17]	313088	57237	San Juan, Vieques, Ponce
Michael	#michael, #hurricanemichael	[06-10-18, 23-10-18]	52506	22007	Panama City, Mexico Beach, Tallahassee

Table 1: Summary statistics for Twitter data.

a diachronic analysis of newspaper articles showed that entities tended to receive fewer descriptors as they became more familiar over time (Staliūnaitė et al. 2018). In laboratory settings, speakers have been found to modulate their use of contextualizing information depending on the expected entity salience to *specific* listeners (Galati and Brennan 2010), but such fine-grained analyses have not previously been explored with corpus data. This study contributes to this literature by extending the focus from professional print media and laboratory experiments to social media, where it is possible to measure fine-grained social and informational factors in naturally-occurring text.

We choose crisis events as the domain of study because they present a natural example of wide-scale information sharing with respect to a shared event (Houston et al. 2015). During crisis events, citizens often turn to social media to share their first-hand experiences (Soden and Palen 2018) and to seek information (Varga et al. 2013). Non-governmental organizations (NGOs), government agencies, and media outlets often rely on social media as a barometer for the experiences of people affected by crises (Imran et al. 2015). These organizations’ relief efforts may be informed by the collective attention directed toward actionable needs of crisis victims (Palen and Hughes 2018). Prior work has proposed a variety of computational methods to extract actionable information from social media for crisis responders (Olteanu, Vieweg, and Castillo 2015; Temnikova, Castillo, and Vieweg 2015), such as locations that need aid.

Our work proposes a fine-grained metric for collective attention to provide insight into how authors frame their writing to address audience expectations during crisis events, which can help organizations understand information needs among the public (Murthy and Gross 2017). We discuss further implications for crisis informatics in § 5.1.

3 Data

Crisis events present a useful case study for the development of collective attention, due to the large volume of online participation and uncertainty among event observers towards the situation (Varga et al. 2013). We chose to study the collective attention changes in public discourse related to hurricanes, due to hurricanes’ lasting economic impact, their broad coverage in the news, and their relevance to specific

Event	Authors	Tweets	LOCATION NEs
Florence	186	17624	29066
Harvey	164	31563	50050
Irma	178	45913	77114
Maria	139	11332	18204
Michael	146	8828	14655

Table 2: Summary statistics for active authors on Twitter.

geographic regions. We collected social media data related to five recent hurricanes. The remainder of this section describes the data collection (§ 3.1), location detection (§ 3.2), and descriptor detection (§ 3.3) for the following datasets:

1. Twitter: 2 million public tweets related to 5 major hurricanes, collected in 2017 and 2018.
2. Facebook: 30,000 posts from 60 public groups related to disaster relief in Hurricane Maria, collected in 2017.

3.1 Collection

Twitter Dataset The Twitter posts were collected using hashtags from five major hurricanes: Hurricane Florence (2018), Hurricane Harvey (2017), Hurricane Irma (2017), Hurricane Maria (2017), and Hurricane Michael (2018). We used hashtags that contained the name of the event in full and shortened form, e.g. #Harvey and #HurricaneHarvey for Hurricane Harvey.

During 2017 and 2018, we streamed tweets that contained hashtags related to the natural disasters at the start of each disaster for up to one week after the dissipation of the hurricane.² We augmented this data with additional tweets available in a 1% Twitter sample that contains the related hashtags, restricting our time frame to one day before the formation of the hurricane and one week after the dissipation of the hurricane. Manual inspection revealed minimal noise generated by the inclusion of the name-only hashtags (e.g., #Harvey).

²Dates are based on estimates from the National Oceanic and Atmospheric Administration (NOAA). For example, estimates for Hurricane Harvey are available at https://www.nhc.noaa.gov/data/tcr/AL092017_Harvey.pdf (accessed January 2019).

Phrase patterns	Dependency types	Example
<u>LOCATION</u> + <u>LOCATION_STATE</u>	n/a	<u>San Juan</u> , <u>PR</u>
<u>LOCATION</u> + [<u>LOCATION_CONTEXT</u>] _{MODIFIER}	adjective, apposition, preposition, numeric modifier	<u>San Juan</u> , [capital of <u>Puerto Rico</u>]
[<u>LOCATION</u> + <u>LOCATION_CONTEXT</u>] _{NOUN_COMPOUND}	nominal, compound, apposition	the [<u>Vega Alta neighborhood of San Juan</u>]
<u>LOCATION</u> + [<u>LOCATION_STATE</u>] _{CONJUNCTION}	conjunction	<u>San Juan</u> , <u>Guayama</u> [and <u>Vieques</u> , <u>Puerto Rico</u>]

Table 3: Phrase patterns to capture descriptor phrases in location mentions. Head location marked with underline, context location marked with double underline.

Summary statistics of the Twitter data are presented in Table 1. We also collected additional event-related tweets from the most frequently-posting authors in each dataset (“active authors”), which were needed to evaluate per-author changes (RQ2b; see § 4.3). Table 2 summarizes the detailed statistics about the active author data.

Facebook Dataset The Facebook data was collected in the aftermath of Hurricane Maria by searching for public discussion groups that included at least one of Puerto Rico’s municipalities in the title (e.g. “*Guayama: Huracán Maria*” refers to Guayama municipality). Relatives and friends of Puerto Ricans often posted in these groups to seek additional information about those still on Puerto Rico, who could not be reached by telephone due to infrastructure damage. We restricted our analysis to Facebook groups related to Hurricane Maria because the limited information available caused more discussion of specific locations, as compared to the other hurricane events that had more up-to-date information available online.

In total, we collected 31,414 public posts from 61 groups, from the time of their creation to one month afterward (September 20 to October 20, 2017). Spanish was the majority language in these posts, so only posts in Spanish were retained, using `langid.py` (Lui and Baldwin 2012).³ Due to Facebook data restrictions and API changes, we were unable to collect posts in Facebook groups for the other four hurricane events.

3.2 Extracting and Filtering Locations

We extracted mentions of locations using two systems for named entity recognition (NER): for English, we used a system that was explicitly adapted to Twitter data (Ritter et al. 2011)⁴ and for Spanish, we used a general purpose named entity recognizer (Finkel, Grenager, and Manning 2005).⁵ These systems are freely accessible and widely used, and achieve reasonably competitive performance.⁶ We manually

³<https://github.com/saffsd/langid.py> (accessed October 2017)

⁴https://github.com/aritter/twitter_nlp (accessed January 2019)

⁵<https://nlp.stanford.edu/software/stanford-ner-2018-10-16.zip> (accessed January 2019)

⁶For location entities, the English tagger has a reported F1 of 74% (Ritter et al. 2011), and the Spanish tagger has a reported F1 of 58% (Finkel and Manning 2009), but these figures are not directly comparable due to genre differences across datasets. Both English and Spanish are considered “high resource” languages for

evaluated the performance of these NER systems on a sample of tweets (100 tagged `LOCATION`s per dataset, 500 total) and found reasonable precision for the `LOCATION` tag (81-96% across all datasets).

For this work, we focus on named entities that may require descriptor phrases, which include cities and counties. We therefore restrict our analysis to named entities (NEs) that (1) are tagged as `LOCATION`, (2) can be found in the GeoNames ontology,⁷ (3) map to cities or counties in the ontology, (4) map to affected locations in the ontology, based on their location occurring in the region affected by the event, and (5) are unambiguous within the region affected by the event. For instance, the string “*San Juan*” is a valid location for the Hurricane Maria tweets because the affected region contains an unambiguous match for the string, but it is not a valid location for the Hurricane Harvey tweets because the affected region does not contain an unambiguous match.

3.3 Extracting Descriptor Phrases

One way in which writer can introduce a new entity to a discourse (e.g., “*San Juan*”) is by linking it to a more well-known entity (e.g., “*Puerto Rico*”) in a descriptor phrase. To detect this phenomenon, we identified location mentions that had dependent clauses that referred to better-known locations, using population as a proxy. The underlying assumption is that a more well-populated location is more likely to be known to readers, and can therefore help describe the preceding location. The frequency of such descriptor phrases is the main dependent variable in this research: we hypothesize that authors are more likely to use such descriptor phrases when they expect readers to treat the location as new information, and less likely to do so when the location is already salient.

To extract sentence structure from text, we used dependency parsing, which decomposes sentences into directed acyclic graphs connecting words and phrases (Eisenstein 2019). Following Staliūnaitė et al. (2018), we defined a set of dependencies to capture the

natural language processing, with hundreds of thousands of tokens of labeled data for named entity recognition (Hovy et al. 2006; Taulé, Martí, and Recasens 2008). The extension of this data acquisition pipeline to languages that lack substantial labeled data may pose a significant challenge (Rahimi, Li, and Cohn 2019).

⁷<http://download.geonames.org/export/dump/allCountries.zip> (accessed September 2017)

Factor	Variable	Description
Importance	Prior location mentions	Frequency of location within the group or event
Author	In-group posts	Posts that an author made within a group
	In-event posts	Posts that an author made about an event (log-transformed)
	In-event posts about location	Posts that an author made about an event that mention the location (log)
	Organization	Whether the author is predicted to be an organization (based on metadata)
	Local	Whether the author is predicted to be local to the event (based on self-reported location)
Audience	Location is local to group	Whether the location exists within the group’s associated region
	Group size	Number of unique members who have posted in the group
	Prior engagement	Mean normalized log-count of retweets and likes received by an author (in t-1)
	Change in prior engagement	Change in prior engagement received by an author (between t-2 and t-1)
Information	Has URL	Whether the post contains a URL
	Has image/video	Whether the post contains a URL with an associated image/video
Time	Time since start	Days since first post about event
	During peak	Whether post was written during peak of collective attention toward location
	Post peak	Whether post was written at least 1 day after the peak of collective attention toward location

Table 4: Summary of explanatory variables and corresponding metrics, used for descriptor phrase prediction.

MODIFIER phrase type in a subclause (adjectival clause, appositional modifier, prepositional modifier, numeric modifier) and another set of dependencies to capture the COMPOUND type in a super-clause (nominal modifier, compound, appositional modifier). Table 3 presents a summary of the phrase patterns that were used to capture descriptor phrases. Taking into account the characteristics of text from two different domains, for the Twitter data we used the `spacy` shift-reduce parser (Honnibal and Johnson 2015)⁸; for the Facebook data, the dependencies were extracted using the SyntaxNet transition-based parser (Andor et al. 2016).⁹ Our pilot experiments found that SyntaxNet achieved higher accuracy on Facebook posts, but we were unable to apply it to the larger Twitter dataset due to API restrictions.¹⁰

Validation of Extraction Performance To assess the accuracy of our phrase patterns in capturing descriptor phrases, we asked two annotators (computer science graduate students) who had not seen the data to annotate a random sample of 50 tweets containing at least one location from each data set (250 tweets total). The annotators received instructions on how to determine if a location

was marked by a descriptor phrase, including examples that were not drawn from the data, and the annotators marked each location mention as either (1) a “LOCATION + LOCATION.STATE” pattern, (2) one of the other descriptor patterns in Table 3 or (3) no descriptor phrase. The annotators achieved high agreement on each separate descriptor type (Cohen’s $\kappa = 0.96$ for the state pattern, $\kappa = 0.91$ for the other patterns). We then filtered posts with perfect agreement, ran dependency parsing on the posts and detected descriptor phrases using the phrase patterns proposed. We found that our phrase patterns achieved reasonable precision and recall (96.6% and 87.5% respectively) in identifying descriptor phrases compared to raters’ annotations. This validation check demonstrated that our proposed syntactic patterns can capture descriptor phrases reasonably well.

4 Results

We address our research questions in three analyses: static social factors, dynamic factors at the collective level, and dynamic factors at the individual level.

4.1 What Affects the Use of Descriptor Phrases?

We first address RQ1, concerning which social factors influence the use of descriptor context when referring to locations of hurricane events. We are particularly interested in indicators of whether locations may be considered shared knowledge within a community. A descriptor phrase may be omitted for locations that are geographically local to a group of people, i.e. knowledge that already shared among

⁸<https://spacy.io/usage> (accessed January 2019)

⁹<https://cloud.google.com/natural-language/docs/analyzing-syntax> (accessed January 2019)

¹⁰As a robustness check we re-ran the analysis for the Facebook data using parses from `spacy`, and found relatively the same effect sizes for all variables considered (see § B in the Appendix).

the group and are therefore assumed to be *old* information (e.g., if someone mentions the location “*San Juan*” in a group based in a region containing San Juan).

We compared the rate of descriptor uses for location mentions in both Facebook and Twitter. For the Facebook data, we determined whether the group’s region contains the location mentioned based on whether the most likely match for the location in the gazetteer is contained in that region.¹¹ We also considered the following additional predictors: frequency with which the location is mentioned in prior posts (importance), author posting frequency in the group (author status), and group size (audience), as summarized in Table 4.

For the Twitter data, we considered the following predictors: location mention frequency in the Twitter sample (importance), whether the author is an organization or a local to the location (author status),¹² whether the post has a URL (information), and whether the post has an image or video (information).

We built separate logistic regression models for the Facebook and Twitter data. In both cases, the dependent variable is whether each location mention was accompanied by a descriptor phrase ($N=18432$ and $N=49020$, respectively). In detail, we used an elastic net regression (Zou and Hastie 2005)¹³ in order to reduce the risk of overfitting. For this analysis, rare categorical values ($N < 20$) for the fixed effects are replaced with RARE values to avoid overfitting to uncommon categories. The columns “*RQ1 (Facebook)*” and “*RQ1 (Twitter)*” in Table 5 report the results of the logistic regression.

On Facebook (see “*RQ 1 (Facebook)*” in Table 5), mentions of locations that are local to the group received significantly fewer contextual descriptors ($\beta = -0.623, p < 0.001$). For example, in the group “*Hurricane Maria in Lajas*” the mention of the municipality “*Lajas*” does not receive an descriptor (“*Do you know if the Bank is open in Lajas?*”), while in the group “*Guayama: Huracán María*” the mention of “*Lajas*” does receive an descriptor (“*People who can bring water to Lajas Puerto Rico: they need water urgently*”).¹⁴ The other predictors did not have a statistically significant effect on descriptor use.

On Twitter, there were several significant effects: more salient and important locations receive fewer descriptors ($\beta = -0.172, p < 0.001$); authors who are local to an event are less likely to include descriptor phrases ($\beta = -0.511, p < 0.001$), while organizational accounts on Twitter are more likely to use descriptor phrases ($\beta = 0.093, p < 0.01$); in posts that contain URLs, descriptors are less likely to appear ($\beta = -0.081, p < 0.05$), but in posts that links to image or video, descriptors are more likely

¹¹When a location string matches multiple location entities, we choose the one with the highest population.

¹²See § A in the Appendix for details on determining whether an author is an organization or local.

¹³An L2 regularization of 0.01 was chosen through grid search to maximize log-likelihood on held-out data (90-10 train/test split).

¹⁴Comments are translated from Spanish and paraphrased to preserve privacy.

($\beta = 0.137, p < 0.001$).

These findings are in accord with the view that authors customize their presentations based on the perceived information needs of readers. Additional context is unnecessary when writing for locals, or when writing about entities that are already salient, or have become salient through repeated mentions. Twitter accounts that represent large organizations are likely writing for large audiences who require more context; locals are more likely writing for their peers, who do not. Additional context can be provided by hyperlinks to detailed stories, but multimedia content such as images and videos do not serve the same purpose, and therefore require additional contextualization.

4.2 Collective Change in Descriptor Context Use

We now turn to a temporal analysis of descriptor use, using longitudinal data from Twitter. As collective attention focuses on affected locations over the course of a crisis event, we those locations to require less contextualization. To test this theory, we augment the predictors from the previous section with two temporal variables: whether the message is posted during or after the peak volume in the discussion of the event, and how many days have elapsed since the start of the hurricane.

The definition of the peak in collective attention is critical, because it determines the point at which an entity is expected to become shared knowledge in a discussion (Staliūnaitė et al. 2018). Following Mitra, Wright, and Gilbert (2016), we defined the time of peak collective attention \hat{t}_i for each location i as the (24-hour) period during which it is mentioned the most frequently: $\hat{t}_i = \arg \max_{t \in T} f_t^{(i)}$, where $f_t^{(i)}$ is the raw frequency of location i at time t (see Figure 1 for peak in “*San Juan*” and “*Myrtle Beach*” mentions). We defined *pre-peak* as the period that ends t_{buffer} days before the frequency peak, *during-peak* as the period at most t_{buffer} days before and at most t_{buffer} days after, and *post-peak* as the period that begins t_{buffer} days after the frequency peak (we set $t_{\text{buffer}} = 1$). As described below, we include fixed effects for authors and locations in robustness checks; to improve stability, we removed all locations that are mentioned on fewer than $N = 5$ separate dates, and combined all authors with only a single post into a RARE bin.

As shown in the “*RQ2a (Twitter)*” column of Table 5, the post-peak time period had less descriptor use than the earlier time periods ($\beta = -0.127, p < 0.001$). Furthermore, descriptor phrase use decreased with the number of days since the start of the event ($\beta = -0.120, p < 0.001$). These findings are consistent with the hypothesis that entities become more salient through the focus of collective attention, and that this salience makes contextualization less necessary. The regression also provides more rigorous validation for the trend shown in Figure 1.

However, an additional potential explanation for the decrease in descriptor context may be a change in the set of authors after the peak in collective attention – for example, an influx of locals, who are less likely to use descriptors overall. To test for this, we re-ran the regression above and

Factor	Variable	RQ1 (Facebook)		RQ1 (Twitter)		RQ2a (Twitter)		RQ2b (Twitter)	
		Estimate	S.E.	Estimate	S.E.	Estimate	S.E.	Estimate	S.E.
Intercept		-2.030	28.550	-1.052*	0.404	-1.026*	0.415	-1.222	11.206
Importance	Prior location mentions	-0.075	7.164	-0.172*	0.025	-0.200*	0.031	-0.107	0.114
Author	Author in-group posts	-0.328	0.522	-	-	-	-	-	-
	Author is organization	-	-	0.093*	0.033	0.092*	0.035	-0.149	0.115
	Author is local	-	-	-0.511*	0.020	-0.797*	0.031	-0.671*	0.107
	Prior event-based posts (from author)	-	-	-	-	-	-	0.110	0.093
	Prior location mentions (from author)	-	-	-	-	-	-	-0.237*	0.091
Audience	Local location	-0.623*	0.106	-	-	-	-	-	-
	Group size	0.121	0.040	-	-	-	-	-	-
	Prior engagement (author)	-	-	-	-	-	-	0.292*	0.052
	Change in prior engagement (author)	-	-	-	-	-	-	-0.004	0.042
Information	Has URL	-	-	-0.081*	0.035	-0.058	0.038	-0.482*	0.154
	Has image/video	-	-	0.137*	0.032	0.124*	0.034	0.562*	0.123
Time	Time since start	-	-	-	-	-0.120*	0.036	-0.004	3.63
	During-peak	-	-	-	-	0.004	0.038	0.144	0.122
	Post-peak	-	-	-	-	-0.127*	0.049	-0.189	0.157
Model deviance			469		2954		4127		2239
Accuracy			71.3%		72.7%		73.3%		75.0%

Table 5: Logistic regression results for all analysis, predicting the presence of a descriptor phrase. All regressions include fixed effects for location, (for Facebook) group, and (for Twitter) event. * indicates $p < 0.05$, otherwise $p > 0.05$ after multiple hypothesis correction. All models have significantly higher log-likelihood as compared to the null model (see “Model deviance”) and significantly higher accuracy than chance (see “Accuracy”; accuracy computed over 10 runs of class-balanced sampled data). Note that the author population for RQ2b is restricted to active authors and therefore different from the author population in RQ1 (Twitter) and RQ2a.

replaced the author variables (“*local*” and “*organization*”) with a fixed effect for each author. We found that the post-peak effect was still significant and negative ($\beta = -0.253, p < 0.05$). This suggests that a change in author population does not explain the decrease in descriptor use over time, or else this would be absorbed by the fixed effects. We note that these findings generally replicate prior work on long-term trends in descriptor phrase usage in non-crisis contexts (Staliūnaitė et al. 2018), although this prior work did not consider attention “peak” as the time variable.

4.3 Individual Change in Descriptor Context Use

We now further examine temporal dynamics at the level of individual authors (RQ2b). Under a strong interpretation of our motivating hypotheses, an author who participates frequently in early discussion of the event may use fewer descriptors later during the event, under the assumption that their readers would no longer need context for their event-related posts. However, other factors may also be at

work: an author who has a growing audience may be more likely to use descriptor phrases to accommodate their new readers.

To better model the author-level changes in descriptor use, we introduced the following additional predictors: number of prior posts by author during event (author-level), number of prior posts by author about the location during event (author-level), engagement received by author at $t - 1$ (audience)¹⁵ and change in engagement received by author between $t - 2$ and $t - 1$ (audience). These predictors required a longitudinal sample of frequently-posting authors, i.e. *active* authors, who were identified as those whose post volumes were at or above the 95th percentile among all authors in our collection. We scraped all publicly available tweets posted by these active authors that mention one of the event’s hashtags during the event time period (e.g.,

¹⁵We define engagement as the mean of retweets and likes, converted into z -scores across the population.

all posts for a Harvey-related active author from between August 17 and September 10, 2017 that use #Harvey or #HurricaneHarvey). The locations and descriptor phrases were processed as described in § 3.2), and we report the relevant statistics for these active authors in Table 2. We built similar regularized logistic regression models only using data from the active authors who posted at least once during each of the time periods, so as to isolate changes for individual authors.

The results are described in the “*RQ2b (Twitter)*” column of Table 5. We find that authors’ prior mentions of a location are associated with less descriptor use ($\beta = -0.237, p < 0.001$) but that there is no significant temporal trend with respect to the start of the event or the peak attention. This latter null result held even when we performed the regression without the additional *author* and *audience* variables. We did find that authors who received more engagement from the audience tended to use more descriptors ($\beta = 0.292, p < 0.001$), which is again consistent with the view that larger audiences necessitate additional contextualization.

We hypothesized that the active authors may be different from the overall population in how they respond to trends in collective attention. To test this, we identified *regular* authors as those with lower post volumes below the 95th percentile, and we re-ran the regression analysis with only these individuals. We found that these less active authors do show a significant decrease in descriptor use following the peak in collective attention ($\beta = -0.127, p < 0.05$) and a decrease in descriptor use over time ($\beta = -0.098, p < 0.05$). It is unclear whether highly active authors have special characteristics, or whether these differences are driven by some other aspect of the design. The set of active authors contains many journalists and news outlets, whose patterns of writing may be shaped by stylistic formalisms but also a greater sensitivity to their audience’s awareness of unfolding situations (Murthy and Gross 2017).

5 Discussion

By examining how people refer to affected locations over the course of crisis events, we found several consistent trends related to audience expectations: (1) When authors are local to a place, or are writing for an audience who is expected to be local, they are less likely to use descriptor phrases to contextualize references to locations, reflecting shared knowledge among the author and audience; (2) At a collective level, descriptor use decreases over time during crisis events, even after controlling for a set of explanatory variables; (3) At an individual level, active authors who receive more prior audience engagement tend to use fewer descriptors, but active authors do not use fewer descriptors over time. This contrasts with the less-active authors who use fewer descriptors over time.

5.1 Implications

Theoretical implications for linguistics The study highlights uses the unique setting of crisis communication to shed new light on how people present information under varying conditions of shared knowledge (Doyle and

Frank 2015). First, when communicating with an audience that is united by geographical affinity, writers tend to omit contextualizing descriptors. This is predicted by the theory of audience design (Bell 1984), which describes how a speaker modifies their language to fit their expected audience. Second, we find a distinction between highly active authors and other participants, which complicates the prior understanding of time as a factor in descriptor use: while less active authors tend to decrease the use of descriptors for an entity over time (replicating the findings of Staliūnaitė et al. 2018), more active authors are insensitive to time, but instead vary their writing style in response to audience reaction. The introduction of these additional predictors gives a more nuanced perspective on how authors anticipate their audience’s information needs. Rather than viewing collective attention as a single quantity among the public, we argue that attention levels vary across communities, and that individual writers make independent judgments about the degree of attention that is likely to be present in their intended audiences.

Implications for crisis monitoring With respect to crisis informatics, this study provides evidence to support theories of information sharing on social media during periods of uncertainty. Local observers are known to share different types of information during crises, as compared to official organizations (Kogan, Palen, and Anderson 2015). We find that local observers tend to use fewer contextualizing descriptors, which is consistent with the idea that they are usually writing for an audience of locals who do not need contextualization. We also find that the use of descriptors decreases after the peak of collective attention during a crisis event. This supports previous studies about information sharing during crisis (Houston et al. 2015), which found that people sharing information during a crisis often spend significant time documenting the details of the crisis. This may mean omitting descriptor phrases (known information) in favor of new details related to the locations affected, which is supported by our finding about URLs.

While we do not offer new tools for crisis informatics, we believe that the findings of this research could help response organizations more effectively track collective attention in crises. Unlike volume-based metrics, our content-based analysis is relatively robust to problems of missing data, which may occur when internet access is lost during a crisis. Since organizations often track emergent needs on social media (Imran et al. 2015), highlighting locations that lose descriptor context over time can help organizations better understand public *awareness*, which may be unavailable from more formal accounts of the crisis. This is especially important in the context of post-crisis recovery (Soden and Palen 2018) in which organizations often need to assess public awareness of locations in need, because a gap in awareness of a location may result in uncertainty among the public. Lastly, this study shows that different groups of people (locals, organizations, active authors) use contextualizing information differently, which can shed light on the situation “on the ground” more effectively.

Identifying authors who consistently omit context for local entities could help response organizations identify people who are more directly involved with the crisis (Yin et al. 2015), who can then be contacted for further information.

5.2 Limitations

Sampling methods We focus on only a set of specific crisis events, chosen mainly due to the large volume of online discussions. It is possible that the patterns observed in our study are specific to these events and locations, so more work is required to establish generalization to other types of crisis events and other types of entities. Even within these events, our data collection relied on hashtags that may not capture the full breadth of discussion of the crisis events, because we may have missed less frequent hashtags that covered other aspects of the discussion. If these hashtags are for some reason unrepresentative, then it is possible that a more extensive dataset might reveal other relationships between descriptor use and information expectations, although we have no reason to believe that this is the case. We focus exclusively on location names because of their geographic relevance to events, but future work should examine other types of named entities (people, organizations) that also undergo change in response to increased attention (Staliūnaitė et al. 2018). Finally, we acknowledge that other online contexts might give rise to different expectations about audience knowledge (e.g., discussion forums for news readers versus online encyclopedias whose text is meant to be relevant long after the crisis has passed), and these platforms may therefore feature different patterns of the use of contextualizing descriptors.

Causation and operationalization Our findings are strictly correlational: for example, we find that post-peak mentions of locations are less likely to contain descriptor phrases, but given our observational setting, we cannot establish that time is a cause of this linguistic difference. We have tried to control for all measurable confounds, but it is always possible that there are unmeasured confounds as well. Second, our operationalization of “peak attention” is based on the number of mentions during a specific crisis event. It is possible that other events attracted attention to the locations under discussion before the crises began, e.g. a political news story relevant to Puerto Rico increased awareness of San Juan before Maria hit, causing a spurious effect. However, we assume that the hurricanes attracted the most attention to the regions of study, given the massive financial and social impact of the crises. Third, our operationalization of descriptor phrases is based on a set of lexico-syntactic patterns, using a natural language processing pipeline. These patterns are not exhaustive, and there may be other syntactic constructions that indicate whether an entity is considered new information for the audience (Rahman and Ng 2011).

5.3 Future work

Future work should investigate more long-term examples of descriptor use change in news media, including cases where descriptors may reemerge after being dropped as in the case of Flint, Michigan (often referred to as “*Flint*” in the early stages of the water crisis that started in 2014). Identifying common trajectories of descriptor use and writing styles across events can provide insight into how public information needs may shift in response to changing crisis conditions (Olteanu, Vieweg, and Castillo 2015). With respect to crises, follow-up work should investigate different types of crisis events to determine whether expectations of shared knowledge are significantly different based on the circumstances (Houston et al. 2015). A fast-moving and highly lethal crisis such as an earthquake may require news media to drop context information quickly to make way for newer or more important information, while a more slow-moving crisis may allow media to retain context to accommodate new readers. Lastly, future work should consider alternate definitions of “context” beyond descriptor phrases. Longer spans of text may include context information in less direct ways (e.g. “*San Juan is flooding. It is the capital of Puerto Rico*”) that may still reflect the author’s assumed need for context.

5.4 Conclusion

This study adds a new content-based perspective to the measurement of collective attention, by analyzing how people discuss breaking news events online. By examining five recent hurricane events on multiple social media platforms, our research demonstrated how the need for contextual information in location entity references is shaped over time by changing expectations of entity salience. This study contributes to a better understanding of how people share information during crisis events, and can extend more broadly to other scenarios that involve wide-scale collective attention.

Acknowledgments

The authors thank the anonymous reviewers, Sandeep Soni and members of Georgia Tech’s SocWeb group for their valuable feedback on earlier drafts. This research was supported by NSF award IIS-1452443 and NIH award R01-GM112697-03. DY is supported in part by a grant from Google.

A Detecting author social status

In the context of event-based public discussions, it is worth considering whether a post author is (1) local and (2) an organization. An author who is *local* (more committed) to the event’s region will already be aware of the locations under discussion (Kogan, Palen, and Anderson 2015) and will be less likely to use context than an author who is unfamiliar with the region’s locations. Organizations such as government agencies are often responsible for disseminating official information to help crisis responders and effectively organize aid (Houston et al. 2015). An author who represents an official organization may want to minimize uncertainty

in their messages and use more context than an author who does not represent an organization, i.e. a citizen observer.

We determine author local status and organization status using a sample of metadata available from archived tweets corresponding to the time periods of interest (covering ~ 20% of all authors in the data). Following prior work in geolocation (e.g. Kariryaa et al. 2018), we approximate the local status of an author posting about an event based on whether the author’s self-reported profile location mentions a relevant city or state in the event’s affected region (e.g. for Hurricane Maria, a local author would mention “*Puerto Rico*” or “*PR*” in their location field).

Factor	Variable	Estimate	S.E.
Intercept		-2.08	34.4
Importance	Prior location mentions	-0.042	8.31
Author	Author in-group posts	-0.172	0.549
Audience	Local location	-0.607*	0.116
	Group size	0.106	40.3
Deviance			469
Accuracy			74.3%

Table 6: Regression results for Facebook data in RQ1, using `spacy` parses to detect descriptor phrases. * indicates $p < 0.05$, otherwise $p > 0.05$ after multiple hypothesis correction.

Organizations are difficult to identify automatically, because there is no single indicator of organization status in a Twitter user’s profile information. To determine whether an author counts as an organization, we apply the pre-trained organization classifier from Wood-Doughty, Mahajan, and Dredze (2018)¹⁶ to the author’s metadata, including name, description, and social attributes.

For both local and organization status, we find reasonable precision with respect to a small subset of hand-labeled authors from our data.¹⁷

B Robustness check for descriptor extraction

As mentioned in §3.3, we used the SyntaxNet parser to extract descriptor phrases from the Facebook data due to the parser’s better performance on longer sentences. To verify the consistency of results across parsers, we re-parsed the Facebook data with the `spacy` parser used for the Twitter and repeated the regression to predict descriptor use from the explanatory factors, i.e. RQ1. The effect sizes and

¹⁶ Accessed 7/2019: <https://bitbucket.org/mdredze/demographer/src/peoples2018/>.

¹⁷ One of the authors annotated 500 accounts as organizations and locals, based on available metadata, and compared these labels to those produced by the local proxy and organization classifier. The local proxy achieved precision of 87% and recall of 58%, and the organization classifier achieved precision of 87% and recall of 54%.

significance remained the same in the regression on `spacy` parses, shown in Table 6 (cf. “*RQ1 (Facebook)*” column in Table 5).

References

- Andor, D.; Alberti, C.; Weiss, D.; Severyn, A.; Presta, A.; Ganchev, K.; Petrov, S.; and Collins, M. 2016. Globally normalized transition-based neural networks. In *ACL*, 2442–2452.
- Bell, A. 1984. Language style as audience design. *Language in society* 13(2):145–204.
- DiJulio, B.; Muñana, C.; and Brodie, M. 2017. Puerto Rico after Hurricane Maria: The Public’s Knowledge and Views of Its Impact and the Response. Technical report, Kaiser Family Foundation.
- Doyle, G., and Frank, M. C. 2015. Shared common ground influences information density in microblog texts. In *NAACL*, 1587–1596.
- Eisenstein, J. 2019. *Introduction to Natural Language Processing*. MIT Press.
- Finkel, J. R., and Manning, C. 2009. Nested named entity recognition. In *EMNLP*, 141–150.
- Finkel, J. R.; Grenager, T.; and Manning, C. 2005. Incorporating non-local information into information extraction systems by Gibbs sampling. In *ACL*, 363–370.
- Galati, A., and Brennan, S. 2010. Attenuating information in spoken communication: For the speaker, or for the addressee? *Journal of Memory and Language* 62(1):35–51.
- Garimella, K.; De Francisci Morales, G.; Gionis, A.; and Mathioudakis, M. 2017. The Effect of Collective Attention on Controversial Debates on Social Media. In *WebSci*, 43–52.
- Grice, H. P. 1975. Logic and conversation. In *Speech acts*. Brill. 41–58.
- Honnibal, M., and Johnson, M. 2015. An improved non-monotonic transition system for dependency parsing. In *EMNLP*, 1373–1378.
- 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.; and Griffith, S. A. 2015. Social media and disasters: A functional framework for social media use in disaster planning, response, and research. *Disasters* 39(1):1–22.
- Hovy, E.; Marcus, M.; Palmer, M.; Ramshaw, L.; and Weischedel, R. 2006. Ontonotes: the 90% solution. In *NAACL*, 57–60.
- Imran, M.; Castillo, C.; Diaz, F.; and Vieweg, S. 2015. Processing Social Media Messages in Mass Emergency: A Survey. *ACM Computing Surveys* 47(4):1–38.
- Kang, J. S.; IV, R. L. L.; Chu, Z.; Chen, Y.; Dua, D.; Gimpel, K.; Singh, S.; and Balasubramanian, N. 2019. PoMo: Generating Entity-Specific Post-Modifiers in Context. In *NAACL*, 826–838.

- Kariryaa, A.; Johnson, I.; Schöning, J.; and Hecht, B. 2018. Defining and Predicting the Localness of Volunteered Geographic Information using Ground Truth Data. In *CHI*.
- Kogan, M.; Palen, L.; and Anderson, K. M. 2015. Think Local, Retweet Global: Retweeting by the Geographically-Vulnerable during Hurricane Sandy. In *CSCW*, 981–993.
- Leavitt, A., and Clark, J. A. 2014. Upvoting Hurricane Sandy: event-based news production processes on a social news site. In *CHI*, 1495–1504.
- Lehmann, J.; Gonçalves, B.; Ramasco, J. J.; and Cattuto, C. 2012. Dynamical Classes of Collective Attention in Twitter. In *WWW*, 251–260.
- Lin, Y.-R.; Keegan, B.; Margolin, D.; and Lazer, D. 2014. Rising Tides or Rising Stars? Dynamics of Shared Attention on Twitter during Media Events. *PLoS ONE* 9(5).
- Lui, M., and Baldwin, T. 2012. langid.py: An off-the-shelf language identification tool. In *ACL*, 25–30.
- Mitra, T.; Wright, G.; and Gilbert, E. 2016. Credibility and Dynamics of Collective Attention. *Proceedings of ACM Human-Computer Interaction* 1:1–17.
- Murthy, D., and Gross, A. J. 2017. Social media processes in disasters: Implications of emergent technology use. *Social Science Research* 63:356–370.
- Olteanu, A.; Vieweg, S.; and Castillo, C. 2015. What to Expect When the Unexpected Happens: Social Media Communications Across Crises. In *CSCW*, 994–1009.
- Palen, L., and Hughes, A. 2018. Social media in disaster communication. In *Handbook of Disaster Research*. 497–518.
- Prince, E. 1992. The ZPG letter: Subjects, definiteness, and information-status. In Mann, W., and Thompson, S., eds., *Discourse Description: Diverse linguistic analyses of a fund-raising text*. 295—325.
- Rahimi, A.; Li, Y.; and Cohn, T. 2019. Massively multilingual transfer for NER. In *ACL*, 151–164.
- Rahman, A., and Ng, V. 2011. Learning the information status of noun phrases in spoken dialogues. In *EMNLP*, 1069–1080.
- Ritter, A.; Clark, S.; Mausam; and Etzioni, O. 2011. Named entity recognition in tweets: an experimental study. In *EMNLP*, 1524–1534.
- Sasahara, K.; Hirata, Y.; Toyoda, M.; Kitsuregawa, M.; and Aihara, K. 2013. Quantifying collective attention from tweet stream. *PLoS One* 8(4):e61823.
- Siddharthan, A.; Nenkova, A.; and McKeown, K. 2011. Information status distinctions and referring expressions: An empirical study of references to people in news summaries. *Computational Linguistics* 37(4):811–842.
- Soden, R., and Palen, L. 2018. Informing crisis: Expanding critical perspectives in crisis informatics. *Proceedings of ACM on Human-Computer Interaction* 2:1–22.
- Staliūnaitė, I.; Rohde, H.; Webber, B.; and Louis, A. 2018. Getting to “Hearer-old”: Charting Referring Expressions Across Time. In *EMNLP*, 4350–4359.
- Taulé, M.; Martí, M. A.; and Recasens, M. 2008. AnCorà: Multilevel Annotated Corpora for Catalan and Spanish. In *LREC*, 96–101.
- Temnikova, I. P.; Castillo, C.; and Vieweg, S. 2015. EMTerms 1.0: A Terminological Resource for Crisis Tweets. In *ISCRAM*, 134–146.
- Varga, I.; Sano, M.; Torisawa, K.; Hashimoto, C.; Ohtake, K.; Kawai, T.; Oh, J.-H.; and De Saeger, S. 2013. Aid is Out There: Looking for Help from Tweets during a Large Scale Disaster. In *ACL*, 1619–1629.
- Wakamiya, S.; Antoine, É.; Jatowt, A.; Kawai, Y.; and Akiyama, T. 2015. Portraying Collective Spatial Attention in Twitter. In *KDD*, 39–48.
- Wood-Doughty, Z.; Mahajan, P.; and Dredze, M. 2018. Johns Hopkins or johnny-hopkins: Classifying Individuals versus Organizations on Twitter. In *Second Workshop on Computational Modeling of People’s Opinions, Personality, and Emotions in Social Media*, 56–61.
- Wu, F., and Huberman, B. A. 2007. Novelty and collective attention. *Proceedings of the National Academy of Sciences* 104(45):17599–17601.
- Yin, J.; Karimi, S.; Lampert, A.; Cameron, M.; Robinson, B.; and Power, R. 2015. Using social media to enhance emergency situation awareness. In *International Joint Conference on Artificial Intelligence*, 4234–4239.
- Zou, H., and Hastie, T. 2005. Regularization and variable selection via the elastic net. *Journal of the royal statistical society: series B (statistical methodology)* 67(2):301–320.