

I Stand With You: Using Emojis to Study Solidarity in Crisis Events

Sashank Santhanam^{1*}, Vidhushini Srinivasan¹, Shaina Glass², Samira Shaikh¹

¹Department of Computer Science, ²Department of Psychology
University of North Carolina at Charlotte
*ssantha1@uncc.edu

Abstract

We study how emojis are used to express solidarity in social media in the context of two major crisis events - a natural disaster, Hurricane Irma in 2017 and terrorist attacks that occurred in November 2015 in Paris. Using annotated corpora, we first train a recurrent neural network model to classify expressions of solidarity in text. Next, we use these expressions of solidarity to characterize human behavior in online social networks, through the temporal and geospatial diffusion of emojis. Our analysis reveals that emojis are a powerful indicator of sociolinguistic behaviors (solidarity) that are exhibited on social media as the crisis events unfold.

1 Introduction

The collective enactment of online behaviors, including prosocial behaviors such as solidarity, has been known to directly affect political mobilization and social movements [Tuf14, Fen08]. Social media, due to its increasingly pervasive nature, permits a sense of immediacy [Gid13] - a notion that produces high degree of identification among politicized citizens of the web, especially in response to crisis events [Fen08]. Furthermore, the multiplicity of views and ideologies that proliferate on Online Social Networks (OSNs) has created a society that is increasingly fragmented and polarized [DVVB⁺16, Sun18]. Prosocial behav-

iors like solidarity then become necessary and essential in overcoming ideological differences and finding common ground [Bau13], especially in the aftermath of crisis events (e.g. natural disasters). Recent social movements with a strong sense of online solidarity have had tangible offline (real-world) consequences, exemplified by movements related to *#BlackLivesMatter*, *#MeToo* and *#NeverAgain* [DCJSW16, Bur18]. There is thus a pressing need to understand how solidarity is expressed online and more importantly, how it drives the convergence of a global public in OSNs.

Furthermore, research has shown that emoticons and emojis are more likely to be used in socio-emotional contexts [DBVG07] and that they may serve to clarify the message structure or reinforce the message content [MO17, DP17]. Riordan [Rio17] found that emojis, especially non-face emojis, can alter the reader's perceived affect of messages. While research has investigated the use of emojis over communities and cultures [BKRS16, LF16] as well as how emoji use mediates close personal relationships [KW15], the systematic study of emojis as indicators of human behaviors in the context of social movements has not been undertaken. We thus seek to understand how emojis are used when people express behaviors online on a global scale and what insights can be gleaned through the use of emojis during crisis events. Our work makes two salient contributions:

- We make available two large-scale corpora¹, annotated for expressions of solidarity using multiple annotators and containing a large number of emojis, surrounding two distinct crisis events that vary in time-scales and type of crisis event.
- A framework and software for analyzing of how emojis are used to express prosocial behaviors such as solidarity in the online context, through the study of temporal and geospatial diffusion of emojis in online social networks.

Copyright © 2018 held by the author(s). Copying permitted for private and academic purposes.

In: S. Wijeratne, E. Kiciman, H. Saggion, A. Sheth (eds.): Proceedings of the 1st International Workshop on Emoji Understanding and Applications in Social Media (Emoji2018), Stanford, CA, USA, 25-JUN-2018, published at <http://ceur-ws.org>

¹https://github.com/sashank06/ICWSM_Emoji

We anticipate that our approach and findings would help advance research in the study of online human behaviors and in the dynamics of online mobilization.

2 Related Work

Defining Solidarity: We start by defining what we mean by solidarity. The concept of solidarity has been defined by scholars in relation to complementary terms such as “community spirit or mutual attachment, social cooperation or charity” [Bay99]. In our work, we use the definition of *expressional solidarity* [Tay15], characterized as individuals expressing empathy and support for a group they are not directly involved in (for example, expressing solidarity for victims of natural disasters or terrorist attacks).

Using Emojis to Understand Human Behavior: With respect to research on expressional solidarity, Herrera et al. found that individuals were more outspoken on social media after a tragic event [HVBGMGS15]. They studied solidarity in tweets spanning geographical areas and several languages relating to a terrorist attack, and found that hashtags evolved over time correlating with a need of individuals to speak out about the event. However, they did not investigate the use of emojis in their analysis.

Extant research on emojis usage has designated three categories, that are 1) function: when emojis replace a conjunction or prepositional word; 2) content: when emojis replace a noun, verb, or adjective; and 3) multimodal: when emojis are used to express an attitude, the topic of the message or communicate a gesture [NPM17]. [NPM17] found that the multimodal category is the most frequently used; and we contend that emojis used in the multimodal function may also be most likely to demonstrate solidarity. Emojis are also widely used to convey sentiment [HGS⁺17], including strengthening expression, adjusting tone, expressing humor, irony, or intimacy, and to describe content, which makes emojis (and emoticons) viable resources for sentiment analysis [JA13, NSSM15, PE15]. We use sentiment of emojis to study the diffusion of emojis across time and region.

To the best of our knowledge, no research to date has described automated models of detecting and classifying solidarity expressions in social media. In addition, research on using such models to further investigate how human behavior, especially a prosocial behavior like solidarity, is communicated through the use of emojis in social media is still nascent. Our work seeks to fill this important research gap.

3 Data Collection

Our analysis is based on social media text surrounding two different crisis attacks: Hurricane Irma in 2017

and terrorist attacks in Paris, November 2015. We begin this section by briefly describing the two corpora.

Irma Corpus: Hurricane Irma was a catastrophic Category 5 hurricane and was one of the strongest hurricanes ever to be formed in the Atlantic². The storm caused massive destruction over the Caribbean islands and Cuba before turning north towards the United States. People took to social media to express their thoughts along with tracking the progress of the storm. To create our Irma corpus, we used Twitter streaming API to collect tweets with mentions of the keyword “irma” starting from the time Irma became an intense storm (September 6th, 2017) and until the storm weakened over Mississippi on September 12th, 2017 resulting in a corpus of >16MM tweets.

Paris Corpus: Attackers carried out suicide bombings and multiple shootings near cafes and the Bataclan theatre in Paris on November 13th, 2015. More than 400 people were injured and over a hundred people died in this event³. As a reaction to this incident, people all over the world took to social media to express their reactions. To create our Paris corpus, we collected >2MM tweets from 13th November, 2015 to 17th November, 2015 containing the word “paris” using the Twitter GNIP service⁴.

Annotation Procedure We performed distance labeling [MBSJ09] by having two trained annotators assign the most frequent hashtags in our corpus with one of three labels (“Solidarity” (e.g. *#solidaritywithparis*, *#westandwithparis*, *#prayersforpuertorico*), “Not Solidarity” (e.g. *#breakingnews*, *#facebook*) and “Unrelated/Cannot Determine” (e.g. *#rebootliberty*, *#syrianrefugees*). Using the hashtags that both annotators agreed upon ($\kappa > 0.65$, which is regarded as an acceptable agreement level) [SW05], we filtered tweets that were annotated with conflicting hashtags from both corpora, as well as retweets and duplicate tweets. Table 1 provides the descriptive statistics of the original (not retweets), non-duplicate tweets, that were annotated as expressing solidarity and not solidarity based on their hashtags that we used for further analysis.

Table 1: Descriptive statistics for crisis event corpora

| # of Tweets | Solidarity | Not Solidarity | Total |
|--------------|------------|----------------|-------|
| Irma | 12000 | 81697 | 93697 |
| Paris | 20465 | 29874 | 50339 |

²<https://tinyurl.com/y843u5kh>

³<https://tinyurl.com/pb2boh>

⁴<https://tinyurl.com/y8amahe6>

4 The Emojis of Solidarity

The main goal of this article is to investigate how individuals use emojis to express a prosocial behavior, in this case, solidarity, during crisis events. Accordingly, we outline our analyses in the form of research questions (RQs) and the resulting observations in the sections below.

RQ1: How useful are emojis as features in classifying expressions of solidarity?

After performing manual annotation of the two corpora, we trained two classifiers for detecting solidarity in text. We applied standard NLP pre-processing techniques of tokenization, removing stopwords and lower-casing the tweets. We also removed hashtags that were annotated from the tweets. Class balancing was used in all models to address the issue of majority class imbalance (count of Solidarity vs. Not Solidarity tweets).

Baseline Models: We used Support Vector Machine (SVM) with a linear kernel and 10 fold cross validation to classify tweets containing solidarity expressions. For the baseline models, we experimented with three variants of features including (a) word bigrams, (b) TF-IDF [MPH08], (c) TF-IDF+Bigrams.

RNN+LSTM Model: We built a Recurrent Neural Network(RNN) model with Long Short-Term Memory (LSTM) [HS97] to classify social media posts into Solidarity and Not Solidarity categories. The embedding layer of the RNN is initialized with pre-trained GloVe embeddings [PSM14] and the network consists of a single LSTM layer. All inputs to the network are padded to uniform length of 100. Table 2 shows the hyperparameters of the RNN model.

Table 3 shows the accuracy of the baseline and RNN+LSTM models in classifying expressions of solidarity from text, where the RNN+LSTM model with emojis significantly outperforms the Linear SVM models in both Irma and Paris corpora.

Table 2: RNN+LSTM model hyperparameters

| Hyperparameters | Value |
|-----------------|-------|
| Batch Size | 25 |
| Learning Rate | 0.001 |
| Epochs | 20 |
| Dropout | 0.5 |

Table 3: Accuracy of the baseline SVM models and RNN+LSTM model

| Accuracy | Irma | Paris |
|---------------------------|--------------|--------------|
| RNN+LSTM (with emojis) | 93.5% | 86.7% |
| RNN+LSTM (without emojis) | 89.8% | 86.1% |
| TF-IDF | 85.71% | 75.72% |
| TF-IDF + Bigrams | 82.62% | 76.98% |
| Bigrams only | 79.86% | 75.24% |









































RQ2: Which emojis are used in expressions of solidarity during crisis events and how do they compare to emojis used in other tweets?






To start delving into the data, in Table 4 we show the total number of emojis in each dataset. We observe that the total number of emojis in the tweets that express solidarity (using ground-truth human annotation) is greater than the emojis in not solidarity tweets, even though the number of not solidarity tweets is greater than the solidarity tweets in both crisis events (c.f. Table 1). We also observe that count of emojis is greater in the Irma corpus than in the Paris corpus, even though the number of solidarity tweets is smaller in the Irma corpus. One reason for this could be that the Hurricane Irma event happened in 2017 when predictive emoji was a feature on platforms, while the Paris event occurred in 2015 when such functionality was not operational.

Table 4: Total number of emojis in each dataset

| # of Emojis | Solidarity | Not Solidarity | Total |
|--------------|------------|----------------|-------|
| Irma | 26197 | 25904 | 52101 |
| Paris | 24801 | 12373 | 37174 |

Table 5: Top ten emojis by frequency and their counts in Irma and Paris corpora

| Rank | Irma Sol. | Irma Not Sol. | Paris Sol. | Paris Not Sol. |
|------|--|--|--|--|
| 1 |  6105 |  2098 |  5376 |  2878 |
| 2 |  2336 |  1827 |  2826 |  1033 |
| 3 |  1977 |  1474 |  2649 |  909 |
| 4 |  1643 |  1193 |  2622 |  779 |
| 5 |  1530 |  823 |  2581 |  760 |
| 6 |  1034 |  794 |  2225 |  616 |
| 7 |  934 |  726 |  1702 |  513 |
| 8 |  820 |  725 |  386 |  510 |
| 9 |  625 |  724 |  340 |  433 |
| 10 |  367 |  683 |  259 |  361 |

To address RQ2, we show in Table 5 the top ten most frequently used emojis across both crisis events in the tweets that express solidarity and those that do not. We observe that  is used more frequently in the Irma solidarity tweets (Rank 3) but not in the Irma tweets that do not express solidarity. In the top 10 Irma emojis used in tweets not expressing solidarity, we also observe more negatively valenced emojis, including  and . The  emoji is interesting, since the prevailing meaning is “face with tears of joy”, however this emoji can sometimes be used to express sadness [WBSD16]. In addition,  is used across all four




Figure 1: Cooccurrence network for emojis expressing solidarity from regions affected by Hurricane Irma

sets, albeit at different ranks (e.g. Rank 1 in Irma solidarity and Rank 6 in Paris solidarity tweets).

When comparing the two crisis events, we make the observation that the top 5 ranked Paris solidarity emojis are flags of different countries, related to expressions of solidarity from these countries, including France (🇫🇷) at Rank 1, while 🇺🇸 appears at Rank 9 in the Irma solidarity set. We can thus observe that even though the underlying behavior we study in these two events is solidarity, the top emojis used to express such behavior are different in the two events. During the Paris event, solidarity is signaled through the use of flag emojis from different countries, while in the Irma corpus flag emojis do not play a prominent role.

RQ3: Which emojis cooccur in tweets that are posted within areas directly affected by crisis events as compared to those tweets that are posted from other areas?

This research question and the two following RQs are driven by the hypothesis that solidarity would be expressed differently by people that are directly affected by the crisis than those who are not [Bue16]. To address RQ3, we first geotagged tweets using geopy Python geocoding library⁵ to map the users' locations to their corresponding country. Table 6 shows the total number of emojis in solidarity tweets that were geotagged and categorized as posted within regions affected by the event vs. other regions. We then built co-occurrence networks of emojis in both Irma and Paris corpora using the R ggnetwork package⁶ with the force-directed layout to compare these emoji co-occurrence networks in solidarity tweets that were posted within areas directly affected by the crisis and the areas that were not (shown in Figures 1-4).

⁵<https://github.com/geopy/geopy>

⁶<https://tinyurl.com/y7xnw9lr>

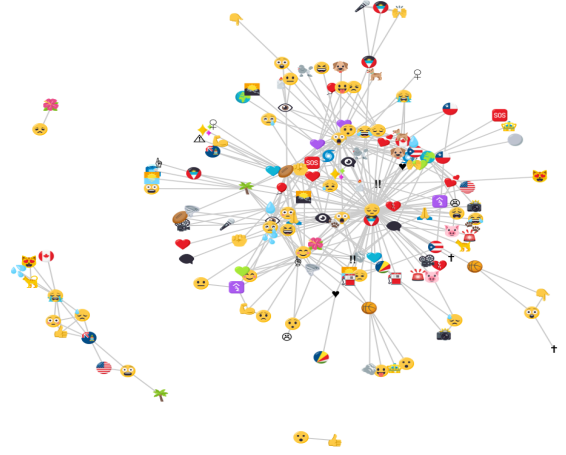


Figure 2: Cooccurrence network for emojis expressing solidarity from regions **not** affected by Hurricane Irma

Table 6: Total count and proportion of emojis in geotagged tweets from affected vs. other regions

| | Irma | Paris |
|------------------|----------------|----------------|
| Affected Regions | 10048 (67.81%) | 925 (6.52%) |
| Other Regions | 4770 (32.19%) | 13267 (93.48%) |

Figure 1 represents the co-occurrence network of emojis within the regions affected by the Hurricane Irma (United States, Antigua and Barbuda, Saint Martin, Saint Barthelemy, Anguilla, Saint Kitts and Nevis, Birgin Islands, Dominican Republic, Puerto Rico, Haiti, Turks and Caicos and Cuba)⁷.

We find the pair 🙏 – ❤️ occurs most frequently in solidarity tweets collected within the Irma affected regions. The other top co-occurring pairs following the sequence include 💙 – 🐷, 🙏 – 🇺🇸, 🇺🇸 – 🇺🇸 and 🇺🇸 – 🙏; these pairs might convey the concerns expressed in the tweets that originate within affected areas. The 💥 emoji appears at the centre of the network denoting the impact of the Irma event. The 🙏 – 🙏, 🙏 – 🙏, 🙏 – 🙏 are the emojis that appear in isolation from the network. The 🙏 and 🙏 emojis can serve as indicators to stay strong during this hurricane calamity.

Figure 2 represents the co-occurrence network of emojis in tweets posted outside the regions affected by Hurricane Irma. We find that the pair 💙 – 🐷 tops the co-occurrence list. Next, we have other co-occurring pairs like 🇺🇸 – 🇺🇸 and 🇺🇸 – 🙏 following the top most frequent pair in sequence. We see the three disjoint networks apart from the main co-occurrence network. The 😞 emoji appears at the centre of the network expressing sorrow and the concern of the people during the event. The disjoint networks also contain flags and other emojis that express sadness and sorrow.

⁷<http://www.bbc.com/news/world-us-canada-41175312>




Figure 3: Cooccurrence network for emojis expressing solidarity from regions affected by November 2015 terrorist attacks in France

Figure 3 shows the co-occurrence network of emojis for November terrorist attacks in France. Within France, the pair 🇫🇷 – ❤️ tops all the co-occurrence pairs. Co-occurrence pairs like 🙏 – 🇫🇷 and 🙏 – ❤️ follow the top co-occurring sequence, strongly conveying the solidarity of people who tweeted during November terrorist attacks. We also have co-occurring pairs containing flags of other countries following the top-tweeted list that shows uniform feeling among the people by trying to express their sorrow and prayers. The 🚫 emoji appears in the centre of the large network as an expression of danger during terrorist attacks. We can also see that the network contains many flags that indicates the concern and worries of people from many different countries. There are five disjoint networks that again contain emojis that express the sorrow, prayers and discontent.

Figure 4 represents the co-occurrence network of emojis for November terrorist attacks in Paris outside France. We find that the pair 🇫🇷 – ❤️ tops the co-occurrence list as within France, which is followed by the co-occurring pairs 🙏 – 🇫🇷 and 🙏 – ❤️. We can infer that the people within or outside France shared common emotions that includes a mixture of prayers, support and concern towards Paris and its people. We find the 🇫🇷, 💖 and 🤝 appear at the centre of the network to convey solidarity. The 🌟 emoji also appears at the center, similar to Figure 1. One important inference is that the 🌟 emoji appears at the network center Irma affected regions whereas it appears at the network center for *unaffected* regions in Paris event.

RQ4: How can emojis be used to understand the diffusion of solidarity expressions over time?

For addressing this research question, we plot in




Figure 4: Cooccurrence network for emojis expressing solidarity from regions **not** affected by November 2015 terrorist attacks in France

Figures 5 and 6 the diffusion of emojis across time (filtering emojis that occur fewer than 50 times and 25 times per day resp. for the 26197 emojis in Irma and 24801 emojis in the Paris solidarity corpus (c.f. Table 4). The emojis are arranged on y-axis based on their sentiment score based on the publicly available work done by Novak et al. [NSSM15].

In the Irma corpus, the temporal diffusion of emojis is quite interesting (Figure 5). Hurricane Irma grabbed attention of the world on September 6th when it turned into a massive storm and the reaction on social media expressing solidarity for Puerto Rico was through ❤️ and 🙏. During the following days, the United States is in the path of the storm, and there is an increased presence of 🇺🇸 and the presence of other countries flags. As the storm lashes out on the islands on September 7th, people express their feelings through ❤️ and 🙏 emojis and also warn people about caring for the pets. As the storm moves through the Atlantic, more prayers with ❤️ and 🙏 emojis emerge on social media for people affected and on the path of this storm. The storm strikes Cuba and part of Bahamas on September 9th before heading towards the Florida coast. As the storm moves towards the US on September 10th, people express their thoughts through 🌊, 🌪️ and 🌪️ causing tornadoes. When images of massive flooding emerge on social media, people respond with pet emojis like 🐱, 🐶, 🐱 and 🐶 to save them. The 🐱 emoji may also serve as an indicator of high-pitched crying [WBSD16].

In the Paris attacks, the first 24 hours from 13th night to 14th were the days on which most number of emojis were used. When news of this horrible attack spreads on social media, the immediate reaction of the people was to express solidarity through hashtags at-



Figure 5: Diffusion of emojis across time for the Hurricane Irma disaster (N=26197 emojis)

tached with 🇫🇷 emoji. As a result, 🇫🇷 was the most frequently used emoji across all days. Emojis of other country flags such as 🇺🇸, 🇬🇧 emerge to indicate solidarity of people from these countries with France. Even after the end of the attack on Nov 13th, people express prayers for the people of France through 🙏 emoji. Images and videos of the attacks emerge on social media on Nov 14th leading to the use of the 📺 emoji. The ❤️ emoji occurs across all days for the Paris event.

Across both events, we also observe a steady presence over all days of positively valenced emojis in the tweets expressing solidarity (the top parts of the diffusion graphs), while negatively valenced emojis are less prevalent over time (e.g. 😡 appears in the first two and three days in the Irma and Paris events resp.).

RQ5: How can emojis be used to study the temporal and geographical diffusion of solidarity expressions during crisis events?

Our primary aim with this research question was to look at how the emojis of solidarity diffuse over time within the affected community and compare communities not affected by the same event. Using the geotagged tweets described in RQ2, we created Figures 7 and 8 to represent the diffusion of emojis over time within the affected regions on the path of Hurricane Irma and the non-affected regions respectively. Since the distribution of emojis in geotagged tweets for the Paris attacks is skewed (6.52% emojis expressing solidarity from affected regions, c.f. Table 6), we have excluded the Paris event for analysis in this RQ.

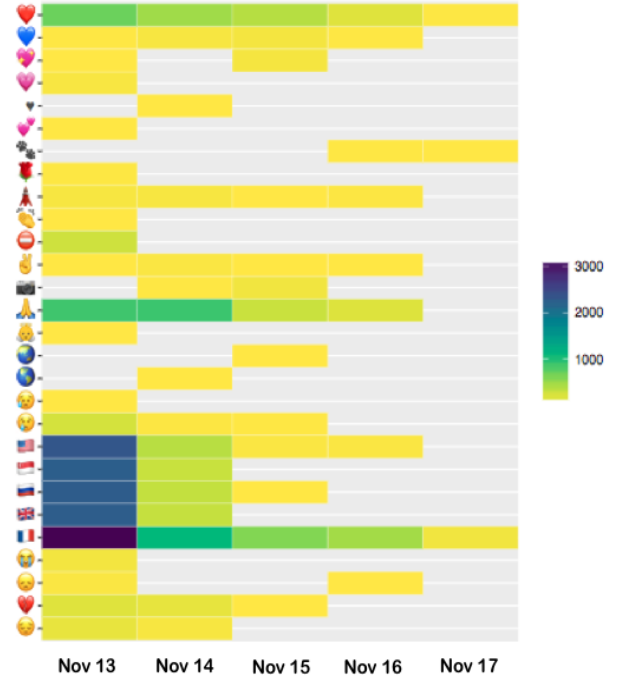


Figure 6: Diffusion of emojis across time for the November Paris attacks (N=24801 emojis)

Figures 7 and 8 allow us to contrast how the Hurricane Irma event is viewed within and outside the affected regions. On September 6th, 🇺🇸 emerged when the hurricane started battering the islands along with a lot of heart emojis. On the other hand, more heart emojis emerged expressing solidarity from outside the affected regions when the people realized the effect of the hurricane (Sep 7th). As the storm moved forward, people in the affected regions express a lot of prayers. On September 9th, as the storm moves towards the United States after striking Cuba, there seems to be more prayers 🙏 amongst affected as well as outside communities. The 🙏 emoji is constant in the affected regions. In both Figures 8 and 7, we see that variety of emojis appear in the latter days of the event (starting Sep 9th), including the 🇨🇦, 😭, and 🐶 as well as the animal/pet emojis, likely indicating the emergence of different topics of discourse related to the event.

5 Conclusion and Future Work

We described our data, algorithm and method to analyze corpora related to two major crisis events, specifically investigating how emojis are used to express solidarity on social media. Using manual annotation based on hashtags, which is a typical approach taken to distance label social media text from Twitter [MBSJ09], we categorized tweets into those that express solidarity and tweets that do not express solidarity. We then analyzed how these tweets and the

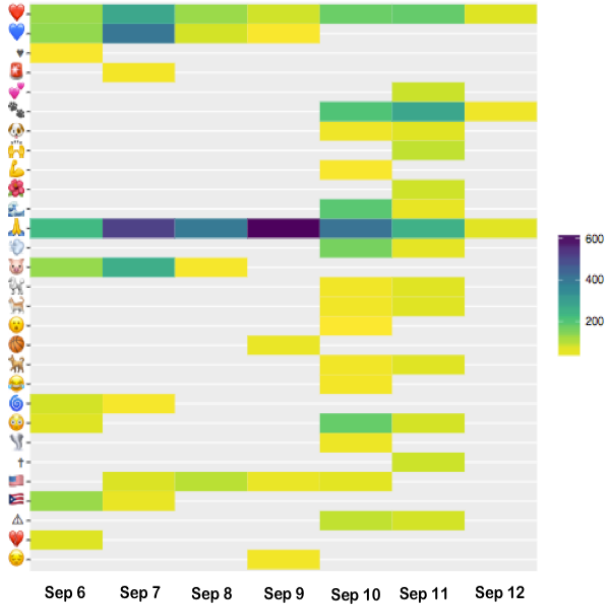


Figure 7: Diffusion of emojis across time for the Hurricane Irma within affected regions (N=10048 emojis)

emojis within them diffused in social media over time and geographical locations to gain insights into how people reacted globally as the crisis events unfolded. We make the following overall observations:

- Emojis are a reliable feature to use in classification algorithms to recognize expressions of solidarity (RQ1).
- The top emojis for the two crisis event reveal the differences in how people perceive these events; in the Paris attack tweets we find a notable presence of flag emojis, likely signaling nationalism but not in the Irma event (RQ2).
- Through the cooccurrence networks, we observe that the emoji pairs in tweets that express solidarity include anthropomorphic emojis (🙏, 😊, ❤️) with other categories of emojis such as 🤝 and 🇫🇷 (RQ3).
- Through analyzing the temporal and geospatial diffusion of emojis in solidarity tweets, we observe a steady presence over all days of positively valenced emojis, while negatively valenced emojis become less prevalent over time (RQ4, RQ5).

Future Work: While this paper addressed five salient research questions related to solidarity and emojis, there are a few limitations. First, our dataset contains emojis that number in the few thousand, which is relatively small when compared to extant research in emoji usage [LF16]. However, we aim to reproduce our findings on larger scale corpora in the




Figure 8: Diffusion of emojis across time for the Hurricane Irma outside affected regions (N=4770 emojis)

future. Second, we analyzed solidarity during crisis events including a terrorist attack and a hurricane, whereas solidarity can be triggered without an overt shocking event, for example the #MeToo movement. In future work, there is great potential for further investigation of emoji diffusion across cultures. In addition to categorizing tweets as being posted from affected regions and outside of affected regions, we wish to analyze the geographical diffusion with more granularity, using countries and regions as our units of analysis to better understand the cultural diffusion of solidarity emojis. An additional future goal is to analyze the interaction of sentiment of emojis and solidarity as well as the text that cooccurs with these emojis in further detail. We anticipate our approach and findings will help foster research in the dynamics of on-line mobilization, especially in the event-specific and behavior-specific usage of emojis on social media.

Acknowledgements

We are grateful for the extremely helpful and constructive feedback given by anonymous reviewers. We thank the two trained annotators for their help in creating our dataset. This research was supported, in part, by a fellowship from the National Consortium of Data Science to the last author.

References

- [Bau13] Zygmunt Bauman. *Postmodernity and its Discontents*. John Wiley & Sons, 2013.

- [Bay99] Kurt Bayertz. Four uses of solidarity. In *Solidarity*, pages 3–28. Springer, 1999.
- [BKRS16] Francesco Barbieri, German Kruszewski, Francesco Ronzano, and Horacio Saggion. How cosmopolitan are emojis?: Exploring emojis usage and meaning over different languages with distributional semantics. In *Proceedings of the 2016 ACM on Multimedia Conference*, pages 531–535. ACM, 2016.
- [Bue16] Steven M Buechler. *Understanding social movements: Theories from the classical era to the present*. Routledge, 2016.
- [Bur18] INQUIRER.net US Bureau. never-again walkouts spark half a million social media posts in 24 hours, Mar 2018.
- [DBVG07] Daantje Derks, Arjan ER Bos, and Jasper Von Grumbkow. Emoticons and social interaction on the internet: the importance of social context. *Computers in human behavior*, 23(1):842–849, 2007.
- [DCJSW16] Munmun De Choudhury, Shagun Jhaver, Benjamin Sugar, and Ingmar Weber. Social media participation in an activist movement for racial equality. In *ICWSM*, pages 92–101, 2016.
- [DP17] Giulia Donato and Patrizia Paggio. Investigating redundancy in emoji use: Study on a twitter based corpus. In *Proceedings of the 8th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 118–126, 2017.
- [DVVB⁺16] Michela Del Vicario, Gianna Vivaldo, Alessandro Bessi, Fabiana Zollo, Antonio Scala, Guido Caldarelli, and Walter Quattrociocchi. Echo chambers: Emotional contagion and group polarization on facebook. *Scientific reports*, 6:37825, 2016.
- [Fen08] Natalie Fenton. Mediating solidarity. *Global Media and Communication*, 4(1):37–57, 2008.
- [Gid13] Anthony Giddens. *The consequences of modernity*. John Wiley & Sons, 2013.
- [HGS⁺17] Tianran Hu, Han Guo, Hao Sun, Thuy-vy Thi Nguyen, and Jiebo Luo. Spice up your chat: The intentions and sentiment effects of using emoji. *arXiv preprint arXiv:1703.02860*, 2017.
- [HS97] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [HVBMGMS15] Enrique Herrera-Viedma, Juan Bernabe-Moreno, Carlos Porcel Gallego, and Maria de los Angeles Martinez Sanchez. Solidarity in social media: when users abandon their comfort zone-the charlie hebdo case. *REVISTA ICONO 14-REVISTA CIENTIFICA DE COMUNICACION Y TECNOLOGIAS*, 13(2):6–22, 2015.
- [JA13] Tanimu Ahmed Jibril and Mardziah Hayati Abdullah. Relevance of emoticons in computer-mediated communication contexts: An overview. *Asian Social Science*, 9(4):201, 2013.
- [KW15] Ryan Kelly and Leon Watts. Characterising the inventive appropriation of emoji as relationally meaningful in mediated close personal relationships. *Experiences of Technology Appropriation: Unanticipated Users, Usage, Circumstances, and Design*, 2015.
- [LF16] Nikola Ljubešić and Darja Fišer. A global analysis of emoji usage. In *Proceedings of the 10th Web as Corpus Workshop*, pages 82–89, 2016.
- [MBSJ09] Mike Mintz, Steven Bills, Rion Snow, and Dan Jurafsky. Distant supervision for relation extraction without labeled data. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP: Volume 2-Volume 2*, pages 1003–1011. ACL, 2009.

- [MO17] Kris M Markman and Sae Os-
hima. Pragmatic play? some pos-
sible functions of english emoticons
and japanese kaomoji in computer-
mediated discourse. 2017.
- [MPH08] Christopher Manning, Raghavan
Prabhakar, and Schütze Hinrich. In-
troduction to information retrieval.
*An Introduction To Information
Retrieval*, 151:177, 2008.
- [NPM17] Noa Na’aman, Hannah Provenza,
and Orion Montoya. Varying linguis-
tic purposes of emoji in (twitter) con-
text. In *Proceedings of ACL 2017,
Student Research Workshop*, pages
136–141, 2017.
- [NSSM15] Petra Kralj Novak, Jasmina
Smailović, Borut Sluban, and
Igor Mozetič. Sentiment of emojis.
PloS one, 10(12):e0144296, 2015.
- [PE15] Umashanthi Pavalanathan and Ja-
cob Eisenstein. Emoticons vs.
emojis on twitter: A causal in-
ference approach. *arXiv preprint
arXiv:1510.08480*, 2015.
- [PSM14] Jeffrey Pennington, Richard Socher,
and Christopher Manning. Glove:
Global vectors for word representa-
tion. In *Proceedings of the 2014 con-
ference on empirical methods in nat-
ural language processing (EMNLP)*,
pages 1532–1543, 2014.
- [Rio17] Monica A Riordan. The communica-
tive role of non-face emojis: Affect
and disambiguation. *Computers in
Human Behavior*, 76:75–86, 2017.
- [Sun18] Cass R Sunstein. *# Republic: Di-
vided democracy in the age of social
media*. Princeton University Press,
2018.
- [SW05] Julius Sim and Chris C Wright. The
kappa statistic in reliability stud-
ies: use, interpretation, and sample
size requirements. *Physical therapy*,
85(3):257–268, 2005.
- [Tay15] Ashley E Taylor. Solidarity: Obl-
igations and expressions. *Journal of
Political Philosophy*, 23(2):128–145,
2015.
- [Tuf14] Zeynep Tufekci. Social move-
ments and governments in the digi-
tal age: Evaluating a complex land-
scape. *Journal of International Af-
fairs*, pages 1–18, 2014.
- [WBSD16] Sanjaya Wijeratne, Lakshika Bala-
suriya, Amit Sheth, and Derek Do-
ran. Emojinet: Building a ma-
chine readable sense inventory for
emoji. In *International Conference
on Social Informatics*, pages 527–
541. Springer, 2016.