Examining the Impact of Social Distance on the Reaction to a Tragedy A Case Study on Sulli's Death

Chen Ling and Gianluca Stringhini
Boston University
ccling@bu.edu, gian@bu.edu

Abstract

In this paper, we aim to gain a better understanding of how social media discussion unfolds in reaction to a tragedy, by focusing on how the social distance between users and the victim impacts them. We leverage tweets regarding the reaction to the death of 25-year old Korean pop star Sulli, who experienced cyber-bullying, depression, and moral coercion by a patriarchal society. We collect 71,588 tweets covering 73 days, characterizing users based on their retweet behavior, and analyzing how they distribute information.

We evaluate the role of official accounts and influential regular accounts in driving the discussion on this tragedy on Twitter and propose a novel multi-language sentiment analysis (English, Korean, Thai) of such discussion based on look-up tables. We then separate users based on their social distance to the victim, based on cultural background (i.e., whether the users come from an East Asian culture) and interest (i.e., whether the accounts are mainly posting about Korean Pop).

Our findings demonstrate that the in-group (i.e., those closer to the tragedy) shows a longer attention period and more frequent in-group interactions compared to the out-group. We notice that the K-pop community is more efficient in spreading information about the tragedy. Our findings describe the information dissemination process after a tragedy and provide insight into potential intervention measures in preventing irrational sensation after a tragedy.

Introduction

People's reaction to tragic events on social media is not well understood by the research community. Previous work focused on crisis events on Twitter, focusing on providing measurements and mitigation methods (Mendoza, Poblete, and Castillo 2010; Cameron et al. 2012; Gruber et al. 2015). A thorough investigation of users' reactions towards tragedy on social media can fill this gap and help us better understanding an Internet-related tragedy (Mendoza, Poblete, and Castillo 2010). In this paper, we take the first step in this direction, by leveraging over 70,000 tweets posted before and after the tragic death of 25-year-old Korean Pop singer Sulli. Based on the social distance theory, we group social network users into in-group or out-group by extracting features reflecting their social distance to the victim. We then analyze how social distance influences the length and the tone of their discussion.

Sulli Choi was a famous South Korean pop star who died on October 14, 2019, at 25 years old. As a young female celebrity, Sulli experienced moral coercion under a patriarchal society in Korea and her rebellion attracted broad attention. Before her death, Sulli had been long harassed by cyberbullying, including hate speech, stalking, and threatening. Her agent revealed her severe depression history too. This tragedy is an example of the situation of many young celebrities, affected by a lack of privacy and forced to live their lives in public (Christofides, Muise, and Desmarais 2012; Fogel and Nehmad 2009). Social media make these problems even more extreme, facilitating toxic behavior like Doxing (Froehlich 2017; Hine et al. 2017; Snyder et al. 2017), cyber-bulling (Hosseinmardi et al. 2015; Tarablus, Heiman, and Olenik-Shemesh 2015; Yao, Chelmis, and Zois 2019), and harassment (Chatzakou et al. 2017a; 2017b). The situation is even more serious with young celebrities, as they attract hate and love at the same time.

Besides facilitating toxic activity, social media amplifies the public reaction to a tragedy, too. The sudden death of an influential celebrity generates heated discussions among her fans. Even people who are not familiar with the pop culture participate in the event to express empathy. Additionally, the attention that the topic attracts can lead to sensationalism and false narratives. Sensational coverage of such tragedy on the media leads to a phenomenon known as Werther effect, where people are inspired by the tragedy and attempt to emulate it (Phillips 1974). Making the event more significant, in Korea, over 20 young musicians, artists, actors, and athletes died accidentally during the past 10 years. The Korean entertainment industry has also been frequently criticized for its high-pressure working environment and its linking with political scandals (Fu and Chan 2013).

In this paper, we focus on people's response to the tragedy of Sulli's death, paying particular attention to their social distance to the event. Social distance reflects the psychological distance between people and others (Trope and Liberman 2010; Bogardus 1933; Akerlof 1997). In psychology, it was shown that this distance affects decision making too.

Research studying social distance is usually performed in two ways. The first method is comparing decisions that subjects make for themselves compared to those they make for others (Polman 2012; Hoffman, McCabe, and Smith 1996). Another method is comparing decisions made for a close friend and a distant friend (or even an unknown person) (Polman and Emich 2011). These manipulation experiments, however, are only conducted in a laboratory environment and have never been tested on observational ground truth data.

In this paper, we aim to bridge this gap by conducting a mixed-method study to build ground truth based on Twitter data according to social distance theory. Specifically, we design an analysis pipeline to study the information diffusion process surrounding a tragedy, focusing on users that are closer to it (in-group), and those that are further away (out-group). We select cultural proximity, and interest in the topic as our parameters of social distance to observe the in-group and out-group difference. Based on these groups, we compare their attention period to the tragedy, their sentiment, and interactions inter/intra groups. While doing so, we propose a way to conduct sentimental analysis for content in multiple languages.

Our contributions include an increased understanding of information distributors in the discussion on Sulli's death and the underlying nature of the narratives towards the tragedy, which may help researchers better understand causes of depression (Lin et al. 2016; Guntuku et al. 2017; Jelenchick, Eickhoff, and Moreno 2013), conspiratorial content (Marwick and Lewis 2017; Starbird et al. 2014; Starbird 2017), Internet campaigns (Altınay 2014; Clark 2016), and cyberbullying (Tippett and Kwak 2012; Calvin et al. 2015; Smith 2013; Aboujaoude et al. 2015). We find evidence that on Twitter, the community that is closer to the tragedy because of interest (i.e., K-pop users) shows more engagement and an discusses the tragedy for a longer period of time, while the same does not happen with the community that is closer due to cultural background (i.e., East Asian accounts). We also find that in the information dissemination process, a higher proportion of in-group members act as an information distributor than the out-group members. In-group members participate in the topic more vividly, acting both as information distributors and retweeters. Members of the K-pop community as in-group also show a tendency of in-group favoritism comparing to the out-group, where negative content on the news of the tragedy is balanced into more neutral tones by the comments of the in-group members.

RELATED WORK

Social media provide a lens for understanding online interactions (Wu et al. 2011; Fischer and Reuber 2011; Kumar et al. 2018; Lazer et al. 2009). Previous work focused on crisis events, leveraging the high volume of real-time data to monitor narratives, as well as to detect the source of conspiratorial content (Starbird 2017; Marwick and Lewis 2017; Starbird et al. 2014). A crisis can be viewed the same as an emergency, i.e. an event or situation that comes on quickly, often without warning, with a potential threat that needs short time response (Ulmer, Sellnow, and Seeger 2017; Acar and Muraki 2011; Schultz, Utz, and Göritz 2011). Twitter works is used to spread real-time news during crisis events, with users often struggling to confirm the reliability of the information they receive (Acar and Muraki 2011; Starbird 2017). Unlike crisis events such as mass shootings, a tragedy arouses more thoughts and rumination instead of fear and nervousness, especially if such tragedy is an uncontrollable event that happens to innocent people (Roland and Munthe 2017).

In front of Sulli's tragedy, social network platforms provide a perfect lens for observing how social distance impact on a victim used to be cyberbullied and depressed based on users' opinions towards the event (Huang et al. 2015). Outside of social media, social distance has been examined as an essential concept in sociology, describing the distance between different groups in society. Social class, race/ethnicity, gender, are the typical categories used in social distance. Social distance measures the intimacy that an individual or group feels towards another individual or group in a social network (Boguná et al. 2004). It can also be used to scale the level of trust of a group towards another, and the extent of the perceived likeness of beliefs (Helfgott and Gunnison 2008). In traditional sociology research, social distance is measured through direct observations of people interacting, questionnaires, speed decision-making tasks, sociograms, etc. (Polman 2012; Hoffman, McCabe, and Smith 1996). Fruitful studies have been done on the impact of social distance on people's speech (Wolfson 1990; Ouellette-Kuntz et al. 2010; Lee and Gibbs 2015).

Leveraging Twitter allows us to study the impact of social distance at a larger scale. In the context of the Internet, strong in-group favoritism may drive users who share the same social identity into irrational toxic behavior, such as cyberbullying (Riek, Mania, and Gaertner 2006; Aboud 2003; Chatzakou et al. 2017b).

METHODOLOGY

Our approach to characterize users' social distance and analyze their behavior on Twitter involves the following steps: (1) data collection, (2) preprocessing, (3) ground truth building, (4) extracting features, (5) analysis of response of in-group and out-group.

Dataset

For the first step, We use Twitter's free streaming API to collect data. We convert all text to lowercase. To better acquire tweets related to Sulli, we filter tweets using Sulli's name and genitive in different languages, including English, Chinese, and Korean. The terms are as follows: “sulli,” “ sulli,” “#sulli,” “sulliś,” “ sulli ,” “雪莉,” “설리,” “최진리.” We gather Twitter data for 73 days from September 1st - 43 days before Sulli's death, to November 13th - 30 days after the news of Sulli's death, on October 14th, 2019. We then exclude the terms (‘sullivan’) from the dataset as it constituted a common false positive in our data.

This collection contains 71,588 tweets in multiple languages, including English (41.8%), Thai (22.6%), Korean

		Positive	Neutral	Negative
Distributors	English	**0.428**	0.355	0.216
	Korean	0.150	0.245	**0.605**
	Thai	0.380	**0.467**	0.153
	Language Avg	0.320	**0.356**	0.325
Retweeters	English	0.358	**0.425**	0.218
	Korean	0.121	0.329	**0.550**
	Thai	0.446	**0.473**	0.081
	Language Avg	0.308	**0.409**	0.283
Official	English	0.25	**0.516**	0.234
	Korean	0.046	0.123	**0.831**
	Thai	**0.474**	0.421	0.105
	Language Avg	0.257	0.352	**0.390**
Overall	English	0.370	**0.416**	0.220
	Korean	0.120	0.310	**0.570**
	Thai	0.441	**0.472**	0.087
	Language Avg	0.310	**0.397**	0.292

Table 1: Sentiment analysis of retweeters, distributors, and official accounts, and all users

Look-up table	Language	Positive	Negative
"Afinn"	English	878	1,598
"KunsentiLex"	Korean	4,868	9,827
Thai sentimental analysis tool-kit	Thai	512	1,219

Table 2: Look up table word list information

(18.2%), Hindi (7.3%), and other languages (10.1%). We count the day of the tragedy released on October 14th, 2019 as day zero. 43 days before her death are marked as -1, -2, ... , -43.

Only 16 out of the 43 days before Sulli's death generated tweets that mentioned the terms above, resulting in a total of 221 tweets. This indicates that Sulli as a Korean pop singer did not attract much attention on Twitter, a western mainstream social network platform. Our following analysis is focused on the data generated since the day of her death (i.e., the zero-day). The number of tweets on the topic reached a peak on the day of Sulli's death, resulting in 49,982 tweets. The discussion of Sulli's death fades quickly. 22 days after the tragic news was released, only 90 tweets mentioned Sulli.

For each tweet, we include the user id, information of creating time, number of favorites, number of the retweet, number of replies, text of the tweet, and the language of the tweet.

For each user who posted a tweet in our dataset, we collect their number of followers, their number of friends, the total number of posted tweets, whether a verified account or not, and the number of tweets that they posted and are contained in our dataset.

Preprocessing

A peculiarity of our dataset is that it contains tweets in several languages. To be able to perform sentiment analysis, we need to come up with a technique able to account for this. To this end, we propose a novel method based on look-up tables (see Figure 1). Our method provides a systematic approach for shedding light on the spreading of a tragedy and the process of shaping narratives. We choose to focus on the top three languages for sentiment analysis: English, Korean, and Thai, which consist of over 82.53% of the tweets in our dataset. To account for different languages, we use Look-up Tables for sentiment analysis across different languages (see Table 2).

For each language, the Look-up Table provides word lists in two sentiment categories: positive, negative. We check each tweet for occurences of the words in our Look-up Table. The sentiment of the content of each tweet is defined as the category with the highest frequency. We label the sentiment of a tweet as negative when either no word in our Look-up Table appears in a tweet or positive and negative words appear in equal numbers.

All the features are examined in two-time scales, a macro-level and a micro-level. At the macro-level, we generate a whole picture of this event on Twitter based on the 31 days sub-dataset since Sulli's death. At the micro-level, we examine each perspective daily.

Categorizing and labeling

First, we manually group users into three roles, retweeter, information distributors, and official accounts to better understand the information diffusion process. These three roles may overlap. There is also a small number of users who do not belong to any role. In this study, we do not look into users who are neither retweeted nor retweet others' posts. The three categories that we study are defined as follows:

- **Retweeters**. Are those accounts that retweeted at least one message.
- **Information distributors**. Are those accounts that posted original messages (not retweets).
- **Official accounts**.Are those accounts verified by Twitter.

Next, we categorize users along two axes to determine their social distance to the tragedy.

Labeling the East Asian cultural sphere. Since our study is cross-language, we group the users whose languages are within the East Asian Culture sphere (i.e., Korean, Japanese, Vietnamese, Chinese, Thai) as in-group. The rationale is that those users are culturally closer to Korea, where the tragedy happened. The rest of the users are considered as out-group.

Labeling the K-pop community. Another dimension across which we evaluate the social distance from the tragedy is whether a user is part of the K-pop community or not. To this end, we fetch the top 200 information distributors (authors of the most-retweeted posts), the top 200 retweeters (users who retweeted the most frequently), and all the official accounts for manual labeling. These users represent over 80% of retweet behavior on this tragic news. Each account is labeled either as a K-pop related account or as an unrelated account. Korea's news account is labeled as a member of the K-pop community. Accounts that are run by individuals who appear not to be a K-pop fan account or to be focused on western musical/news account are labeled as unrelated. Suspended accounts are labeled as unrelated.

A typical K-pop related account should include any of the following elements:

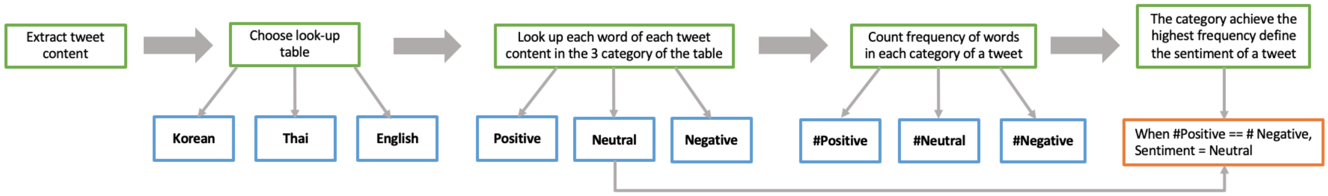

Figure 1: Sentiment Analysis process

- Profile image or cover: a close up shot of a Korean pop singer;
- Profile signature: contains Korean pop singer's name, usually before/after a "_";
- Tweet frequency: tweets actively on Korean pop;
- Tweet content: tweets are all about Korean pop.

These K-pop community members are fans of Korean pop singers or bands. We discard users from fan communities of Korean actors or actresses.

FEATURE EXTRACTION

Our analysis focuses on the users who shared tweets on Sulli's death since the zero-day. 63,797 users generated 71,367 tweets for 30 days from October 14, 2019 to November 13, 2019. 92.7% of users acted as an information distributor or a retweeter in spreading the news of this tragedy. Among all the users, over 80% (53,396 out of 63,992) are retweeters, while 10% (6,317 out of 63,992) are information distributors. Less than 1% of the users take both roles. Only 199 official accounts joined in the discussion. However, nearly 90% of them (177 out of 199) acted as information distributors.

In this section, we describe how each of the different user roles operated while posting about the tragedy, focusing on user quality, tweet features, and tweet content.

User Characteristics

A Chi-square test shows that the profile characteristics of users (number of followers, number of friends, number of total tweets, number of tweets about Sulli) classified as retweeters, information distributors, and official accounts are significantly different ($p < 0.01$). Retweeters have a lower number of followers among these three groups. However, they do follow others. Information distributors have significantly more followers than retweeters and a lower number in total tweets. Interestingly, most of the information distributors do not follow any account on Twitter as the mode is 0. Official accounts have the highest number of followers, friends, and total tweet numbers. Official accounts also post more tweets during the period on Sulli's death.

Tweet Features

Tweet features, including the number of favorite tweets, replies, and retweets, are significantly different among retweeters, information distributors, and official accounts ($p < 0.01$ according to a Chi-square test). Official accounts gain the most favorites and replies. However, the posts by official accounts attract less attention than other information distributors' post after the zero-day. This may imply that most of the official accounts post news on Twitter. After the news has been disseminated, the posts of official accounts about it are not attractive anymore. Although fewer retweets appeared on Twitter after 24 hours of the tragic news, more and more retweeters' posts were being retweeted.

Group	Total	# Official	# Retweeter	# Distributor
CB-In group				
Tweets	29,898	94(0.314%)	26,863(89.84%)	2,234(7.47%)
Users	26,844	36(0.13%)	24,355(90.73%)	1,596(5.95%)
CB-Out group				
Tweets	41,690	368(0.883%)	32,438(77.81%)	5,415(12.99%)
Users	37,617	163(0.43%)	29,481(78.37%)	4,055(10.78%)
IC-In group				
Tweets	758	135(17.81%)	402(53.03%)	431(56.86%)
Users	185	8(4.32%)	91(49.19%)	105(56.76%)
IC-Out group				
Tweets	70,830	327(0.46%)	58,899(83.16%)	7,217(10.19%)
Users	63,809	191(0.29%)	53,305(83.54%)	5,496(8.61%)

Table 3: Tweets and users description statistics according to Cultural Background(CB) and Interest Community(IC)

Tweet Content

We analyze the sentiment of the tweet content posted in English, Korean, and Thai. We assign a general sentiment (positive, negative, or neutral) to each role based on the average score of the tweets posted in each of the three languages (see Table 1). We find that only official accounts present a negative sentiment in their tweet content. Because most official accounts are information distributors, a different sentiment in this group might indicate that information distributors are distributing the tragic news in a different tone of narratives. We assume that official accounts function as public relations, and are therefore more objective in spreading the news of this tragic events.

RESPONSE OF THE IN GROUP & OUT GROUP

We next aim to understand how social distance influences the Twitter discussion on the tragedy. We establish in-group and out-group along two dimensions: national cultural background and community of interest. For the first set, we compare in-group user (Chinese, Japanese, Korean, Thai, Vietnamese language users) and out-group users (the users from other languages). For the second set, we compare in-group users (members of the K-pop community) and out-group users (other users).

National Cultural Background

We find 26,844 users in the in-group 29,898 users in the out-group. A higher percentage of in-group members participate in the information dissemination process. They either act as information distributors, or retweeters: 96.68% of active roles are taken among the in-group members. For the out-group, on the other hand, information distributors and retweets consist of 89.15% of the members. Further examination on the retweeting behavior inter- and intra-groups shows that in-group members interact more frequently, but the result is neither significant nor presents a clear sentiment inclination.

Table 3 shows the number of tweets posted on Sulli's death by the in and out-group. As can be seen, the user activity appears very similar. We then look at the attention period of users in the in- and out-group, defined as the distance between the first and the last tweet that they posted on the topic. We find that in-group users have a longer attention period for this event. The P-value is significant at 0.001 level with D = 0.0958. In other words, users in the in-group discuss the tragedy for longer, compared to users in the out-group.

Interest Community

We label 185 K-Pop community members from a total number of 599 users, including the top 200 retweeting users from the retweeters and the top 200 retweeted users from the information distributors, as well as all 199 official accounts. Please note that there are overlaps of users between these three roles. Considering these 185 members as in-group, the out-group is much larger (63,809 users). Noticeably, the in group shows much enthusiasm following the event. Although there are only 185 members in the in-group, 105 of them (56.76%) act as information distributors and 91 of them (49.19%) are retweeters. The high percentage of both roles indicates an overlap between these two populations. In other words, at least 10% of the posts by in-group users are being retweeted, and these users are also retweeting from others. For the out-group members, the percentage of retweeters is higher than for the in-group: 83.54% of out-group members retweet posts about Sulli. However, the percentage of the information distributors is low: only 8.61% of out-group members' posts have been retweeted.

We then look at the activity of the accounts in the in- and out-group. Compared to the cultural background case, we find that K-Pop accounts discuss the tragedy much more thoroughly. As Figure 2a shows, 30% of the accounts post about Sulli's death 3 or more times, while 90% of the out-group accounts post about it only once. The KS test shows the largest gap between the two groups, D = 0.617, and the P-value of 0 confirms that the two distribution are statistically signifcantly different. When looking at the attention period between the members of the groups, this is also significantly different (see Figure 2b). More in-group members follow the topic of Sulli for several days. Some of the members may tweet on Sulli before her death news releases. We define these members as Sulli's fans. In this case, the P-value is significant at the 0.0001 level with D = 0.21378.

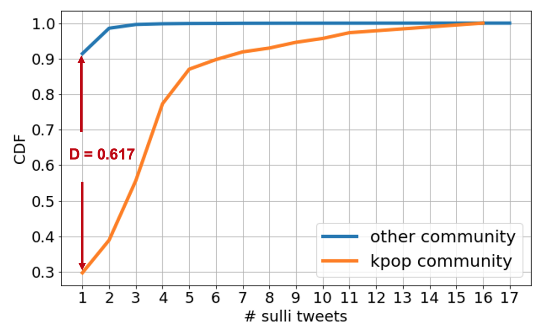

(a) Number of Sulli's related tweets distribution

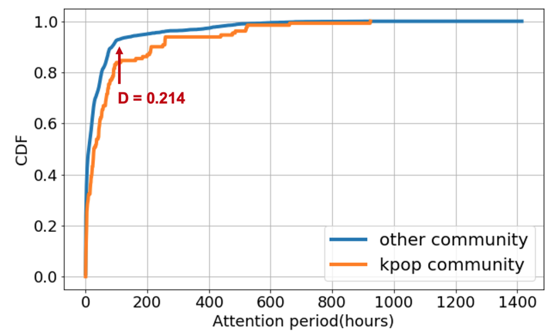

(b) Attention period distribution

Figure 2: Tweets number related to Sulli and attention hours distribution of in-group and out-group between Interest Communities

Figure 3 depicts the interactions between users in the in and out-group. Users in the in-group interact with each other 128 times (among 183 members), with less 0.7 intra-group retweets per member. However, in-group members interact with the out-group members more frequently, with 235 retweets (1.28 per member). The per capita interaction frequency is even lower with the out-group members. The out-group members interact with each other 38,925 times (among 63,809 members), resulting in an average of 0.61 intra-group interaction per member. Inter-group interaction from out-group members to the in-group is as low as 0.31 times per member, with 19,934 retweeted posts from in-group members.

Most of the information flow from the in-group is due to a large number of retweets. Looking closer at the shared content, this information is mostly news articles, and is heavily negative, describing the final scene of Sulli's life. A closer look at the sentiment flow from the in-group, and the interaction between the in-group and the out-group is negative. We assume that this negative sentiment is due to the content of the tragic news.

Conversely, the sentiment that flows from the out-group to the in-group, and the intra-interactions among in-group members, are neutral. We assume that although the news spreading among users is the same, the messages posted when retweeting the news are different between the in-group and the out-group, and are less negative for the in-group.

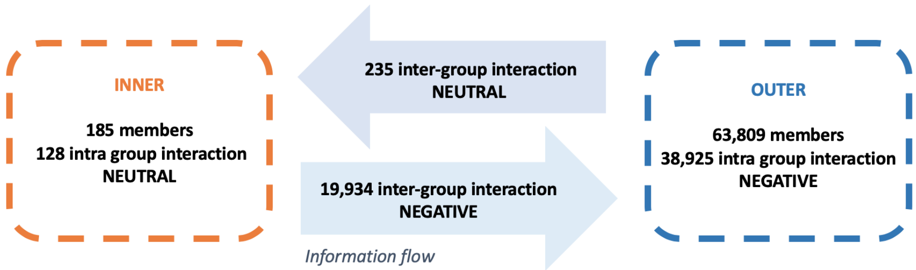

Figure 3: Inter- and Intra- retweet frequency and sentiment between in-group and out-group

DISCUSSION & CONCLUSION

In this paper, we extract ground truth data through the lens of Twitter, to examine how social distance impacts people's response to a tragedy, the sudden passing of a 25-year old K-pop singer. We design a mixed method to process the data, tracing the retweeting behavior to follow the information dissemination on Twitter.

Our results show that when discussing the tragedy, retweeters were the majority of the users. Official accounts represent an extremely small portion of users, however, they act as information distributors. They spread news other than official news, providing attractive narratives or information for the retweeters.

We then investigate how social distance influences user activity on the topic on Twitter. We find that national cultural background may not work as a strong indicator of users' response towards a tragedy: closer cultural and geographic locations do not attract peoples' attention significantly longer periods. On the other hand, we find that similar interest is a strong indicator of users' responses. Considering users in the K-pop community as in-group shows a significantly longer attention period and frequent tweet behavior on the tragic event. These users act both as information distributors and retweeters during the information dissemination. This result is consistent with psychological studies on social distance (Oppenheimer and Olivola 2011; Huang et al. 2015), which found that people tend to care more about those who they know well, or whom they can relate to. In this case, members of the K-pop community knew Sulli better than the other users.

The out-group members are mostly retweeters. At the same time, a higher frequency of intra-group interaction in the out-group aligns with previous studies that social distance affects what kinds of information people learn from other people (Kalkstein et al. 2016).

A slight change (see Figure 3) in sentiment from negative to neutral of the in-group's intra-group interactions and the interactions between the in- and out-group shows a certain level of in-group favoritism (Fu et al. 2012; Balliet, Wu, and De Dreu 2014), too. On the day in which the news of Sulli's death was announced, the "Love you Sulli" campaign was launched on social network platforms. The goal of this campaign was to post positive content to cover the negative information about the tragedy, especially reports of Sulli being cyberbullied, recording of Sulli's live shows under depression, as well as alternative narratives on her death. The campaign was initiated by Sulli's fans, and users in the K-pop community, Korean internet users, and other people who felt empathy towards her heavily participated in it.

Limitations. As all mixed-method approaches, our study has limitations. We trace the distribution of information on twitter based on user retweet behavior, which covers 87% of the tweets of this event. We do not examine the remaining of the 13% tweets, which may include original opinions that did not get re-shared by other users. Limited to the tools to meet language diversity, we do not propose a deeper analysis of the content of the tweets. A further comparison of the NLP tool's impact on the result and a methodology bias should be measured. Different narratives towards the event should be examined based on the ground truth data provided by social network platforms.

In future work, we plan to study the similarities and differences between the reaction to Sulli's death and the one to other tragic events.

References

Aboud, F. E. 2003. The formation of in-group favoritism and out-group prejudice in young children: Are they distinct attitudes? *Developmental psychology* 39(1):48.

Aboujaoude, E.; Savage, M. W.; Starcevic, V.; and Salame, W. O. 2015. Cyberbullying: Review of an old problem gone viral. *Journal of adolescent health* 57(1):10–18.

Acar, A., and Muraki, Y. 2011. Twitter for crisis communication: lessons learned from japan's tsunami disaster. *International Journal of Web Based Communities* 7(3):392–402.

Akerlof, G. A. 1997. Social distance and social decisions. *Econometrica: Journal of the Econometric Society* 1005–1027.

Altınay, R. E. 2014. "there is a massacre of women": violence against women, feminist activism, and hashtags in turkey. *Feminist Media Studies* 14(6):1102–1103.

Balliet, D.; Wu, J.; and De Dreu, C. K. 2014. Ingroup favoritism in cooperation: A meta-analysis. *Psychological bulletin* 140(6):1556.

Bogardus, E. S. 1933. A social distance scale. *Sociology & Social Research.*

Boguná, M.; Pastor-Satorras, R.; Díaz-Guilera, A.; and Arenas, A. 2004. Models of social networks based on social distance attachment. *Physical review E* 70(5):056122.

Calvin, A. J.; Bellmore, A.; Xu, J.-M.; and Zhu, X. 2015. # bully: Uses of hashtags in posts about bullying on twitter. *Journal of School Violence* 14(1):133–153.

Cameron, M. A.; Power, R.; Robinson, B.; and Yin, J. 2012. Emergency situation awareness from twitter for crisis management. In *Proceedings of the 21st International Conference on World Wide Web*, 695–698.

Chatzakou, D.; Kourtellis, N.; Blackburn, J.; De Cristofaro, E.; Stringhini, G.; and Vakali, A. 2017a. Hate is not binary: Studying abusive behavior of# gamergate on twitter. In *Proceedings of the 28th ACM conference on hypertext and social media*, 65–74.

Chatzakou, D.; Kourtellis, N.; Blackburn, J.; De Cristofaro, E.; Stringhini, G.; and Vakali, A. 2017b. Mean birds: Detecting aggression and bullying on twitter. In *Proceedings of the 2017 ACM on web science conference*, 13–22.

Christofides, E.; Muise, A.; and Desmarais, S. 2012. Hey mom, what's on your facebook? comparing facebook disclosure and privacy in adolescents and adults. *Social Psychological and Personality Science* 3(1):48–54.

Clark, R. 2016. "hope in a hashtag": the discursive activism of# whyistayed. *Feminist Media Studies* 16(5):788–804.

Fischer, E., and Reuber, A. R. 2011. Social interaction via new social media:(how) can interactions on twitter affect effectual thinking and behavior? *Journal of business venturing* 26(1):1–18.

Fogel, J., and Nehmad, E. 2009. Internet social network communities: Risk taking, trust, and privacy concerns. *Computers in human behavior* 25(1):153–160.

Froehlich, T. J. 2017. A not-so-brief account of current information ethics: The ethics of ignorance, missing information, misinformation, disinformation and other forms of deception or incompetence. *BiD* (39).

Fu, K.-w., and Chan, C. 2013. A study of the impact of thirteen celebrity suicides on subsequent suicide rates in south korea from 2005 to 2009. *PloS one* 8(1).

Fu, F.; Tarnita, C. E.; Christakis, N. A.; Wang, L.; Rand, D. G.; and Nowak, M. A. 2012. Evolution of in-group favoritism. *Scientific reports* 2:460.

Gruber, D. A.; Smerek, R. E.; Thomas-Hunt, M. C.; and James, E. H. 2015. The real-time power of twitter: Crisis management and leadership in an age of social media. *Business Horizons* 58(2):163–172.

Guntuku, S. C.; Yaden, D. B.; Kern, M. L.; Ungar, L. H.; and Eichstaedt, J. C. 2017. Detecting depression and mental illness on social media: an integrative review. *Current Opinion in Behavioral Sciences* 18:43–49.

Helfgott, J. B., and Gunnison, E. 2008. The influence of social distance on community corrections officer perceptions of offender reentry needs. *Fed. Probation* 72:2.

Hine, G. E.; Onaolapo, J.; De Cristofaro, E.; Kourtellis, N.; Leontiadis, I.; Samaras, R.; Stringhini, G.; and Blackburn, J. 2017. Kek, cucks, and god emperor trump: A measurement study of 4chan's politically incorrect forum and its effects on the web. In *Eleventh International AAAI Conference on Web and Social Media*.

Hoffman, E.; McCabe, K.; and Smith, V. L. 1996. Social distance and other-regarding behavior in dictator games. *The American economic review* 86(3):653–660.

Hosseinmardi, H.; Mattson, S. A.; Rafiq, R. I.; Han, R.; Lv, Q.; and Mishra, S. 2015. Detection of cyberbullying incidents on the instagram social network. *arXiv preprint arXiv:1503.03909*.

Huang, Y. L.; Starbird, K.; Orand, M.; Stanek, S. A.; and Pedersen, H. T. 2015. Connected through crisis: Emotional proximity and the spread of misinformation online. In *Proceedings of the 18th ACM conference on computer supported cooperative work & social computing*, 969–980.

Jelenchick, L. A.; Eickhoff, J. C.; and Moreno, M. A. 2013. "facebook depression?" social networking site use and depression in older adolescents. *Journal of Adolescent Health* 52(1):128–130.

Kalkstein, D. A.; Kleiman, T.; Wakslak, C. J.; Liberman, N.; and Trope, Y. 2016. Social learning across psychological distance. *Journal of personality and social psychology* 110(1):1.

Kumar, S.; Hamilton, W. L.; Leskovec, J.; and Jurafsky, D. 2018. Community interaction and conflict on the web. In *Proceedings of the 2018 World Wide Web Conference*, 933–943.

Lazer, D.; Pentland, A.; Adamic, L.; Aral, S.; Barabási, A.-L.; Brewer, D.; Christakis, N.; Contractor, N.; Fowler, J.; Gutmann, M.; et al. 2009. Computational social science. *Science* 323(5915):721–723.

Lee, J., and Gibbs, J. C. 2015. Race and attitudes toward police: The mediating effect of social distance. *Policing: An International Journal of Police Strategies & Management*.

Lin, L. Y.; Sidani, J. E.; Shensa, A.; Radovic, A.; Miller, E.; Colditz, J. B.; Hoffman, B. L.; Giles, L. M.; and Primack, B. A. 2016. Association between social media use and depression among us young adults. *Depression and anxiety* 33(4):323–331.

Marwick, A., and Lewis, R. 2017. Media manipulation and disinformation online. *New York: Data & Society Research Institute*.

Mendoza, M.; Poblete, B.; and Castillo, C. 2010. Twitter under crisis: Can we trust what we rt? In *Proceedings of the first workshop on social media analytics*, 71–79. ACM.

Oppenheimer, D. M., and Olivola, C. Y. 2011. *The science of giving: Experimental approaches to the study of charity*. Psychology Press.

Ouellette-Kuntz, H.; Burge, P.; Brown, H. K.; and Arsenault, E. 2010. Public attitudes towards individuals with intellectual disabilities as measured by the concept of social distance. *Journal of Applied Research in Intellectual Disabilities* 23(2):132–142.

Phillips, D. P. 1974. The influence of suggestion on suicide: Substantive and theoretical implications of the werther effect. *American Sociological Review* 340–354.

Polman, E., and Emich, K. J. 2011. Decisions for others are more creative than decisions for the self. *Personality and Social Psychology Bulletin* 37(4):492–501.

Polman, E. 2012. Effects of self–other decision making on regulatory focus and choice overload. *Journal of personality and social psychology* 102(5):980.

Riek, B. M.; Mania, E. W.; and Gaertner, S. L. 2006. Intergroup threat and outgroup attitudes: A meta-analytic review. *Personality and social psychology review* 10(4):336–353.

Roland, E., and Munthe, E. 2017. *Bullying (1989): An International Perspective*. Routledge.

Schultz, F.; Utz, S.; and Göritz, A. 2011. Is the medium the message? perceptions of and reactions to crisis communication via twitter, blogs and traditional media. *Public relations review* 37(1):20–27.

Smith, P. K. 2013. School bullying. *Sociologia, problemas e práticas* (71):81–98.

Snyder, P.; Doerfler, P.; Kanich, C.; and McCoy, D. 2017. Fifteen minutes of unwanted fame: Detecting and characterizing doxing. In *Internet Measurement Conference*, 432–444.

Starbird, K.; Maddock, J.; Orand, M.; Achterman, P.; and Mason, R. M. 2014. Rumors, false flags, and digital vigilantes: Misinformation on twitter after the 2013 boston marathon bombing. *IConference 2014 Proceedings*.

Starbird, K. 2017. Examining the alternative media ecosystem through the production of alternative narratives of mass shooting events on twitter. In *Eleventh International AAAI Conference on Web and Social Media*.

Tarablus, T.; Heiman, T.; and Olenik-Shemesh, D. 2015. Cyber bullying among teenagers in israel: An examination of cyber bullying, traditional bullying, and socioemotional functioning. *Journal of Aggression, Maltreatment & Trauma* 24(6):707–720.

Tippett, N., and Kwak, K. 2012. Cyberbullying in south korea. *Cyberbullying in the global playground: Research from international perspectives* 202–219.

Trope, Y., and Liberman, N. 2010. Construal-level theory of psychological distance. *Psychological review* 117(2):440.

Ulmer, R. R.; Sellnow, T. L.; and Seeger, M. W. 2017. *Effective crisis communication: Moving from crisis to opportunity*. Sage Publications.

Wolfson, N. 1990. The bulge: A theory of speech behavior and social distance. *Penn Working Papers in Educational Linguistics* 2(1):55–83.

Wu, S.; Hofman, J. M.; Mason, W. A.; and Watts, D. J. 2011. Who says what to whom on twitter. In *Proceedings of the 20th international conference on World wide web*, 705–714. ACM.

Yao, M.; Chelmis, C.; and Zois, D. S. 2019. Cyberbullying ends here: Towards robust detection of cyberbullying in social media. In *The World Wide Web Conference*, 3427–3433.