

# Contrasting Misinformation and Real-Information Dissemination Network Structures on Social Media During a Health Emergency


Lida Safarnejad, PhD, Qian Xu, PhD, Yaorong Ge, PhD, Siddharth Krishnan, PhD, Arunkumar Bagarvathi, PhD, and Shi Chen, PhD

**Objectives.** To provide a comprehensive workflow to identify top influential health misinformation about Zika on Twitter in 2016, reconstruct information dissemination networks of retweeting, contrast mis- from real information on various metrics, and investigate how Zika misinformation proliferated on social media during the Zika epidemic.

**Methods.** We systematically reviewed the top 5000 English-language Zika tweets, established an evidence-based definition of “misinformation,” identified misinformation tweets, and matched a comparable group of real-information tweets. We developed an algorithm to reconstruct retweeting networks for 266 misinformation and 458 comparable real-information tweets. We computed and compared 9 network metrics characterizing network structure across various levels between the 2 groups.

**Results.** There were statistically significant differences in all 9 network metrics between real and misinformation groups. Misinformation network structures were generally more sophisticated than those in the real-information group. There was substantial within-group variability, too.

**Conclusions.** Dissemination networks of Zika misinformation differed substantially from real information on Twitter, indicating that misinformation utilized distinct dissemination mechanisms from real information. Our study will lead to a more holistic understanding of health misinformation challenges on social media. (*Am J Public Health*. 2020;110:S340–S347. <https://doi.org/10.2105/AJPH.2020.305854>)

 **See also Chou and Gaysynsky, p. S270.**

**S**ocial media have become real-time sources of information on various fields, including health and medical-related topics.<sup>1,2</sup> Contents on a social media platform, such as Twitter, are mainly user-generated, and the lack of effective fact-checking mechanisms makes social media susceptible to propagations of misinformation. Infiltration and proliferation of health-related misinformation on social media, especially during health emergencies, is a serious threat to people and the entire society.<sup>3</sup> Misinformation about vaccines,<sup>4–6</sup> Zika,<sup>7</sup> tobacco, vaping, and marijuana products<sup>8</sup> are a few examples that demonstrate the health-related misinformation problem on social media.

While social media can be an effective tool to enhance people’s health literacy,<sup>9,10</sup> they are also a rich resource to study the public’s

perspectives and reactions toward various topics.<sup>11–14</sup> Infosurveillance systems aim to strengthen the capacity of the public health community by closely monitoring online discussions of health topics<sup>15–18</sup> and detecting misinformation.<sup>19</sup> State-of-the-art analyses of (mis)information dissemination on social media mainly seek 2 purposes: (1) analyzing the information cascade<sup>20,21</sup> and (2)

identifying misinformation.<sup>22–26</sup> In the first direction, computational modeling is used to investigate the virality and spread of (mis)information. In the second direction, different attributes, mainly context-based, are examined to identify misinformation.<sup>23,24,27–33</sup> Nevertheless, developing such systems based on textual content is challenging.<sup>34</sup> The content can be altered to appear real to avoid being detected by automated algorithms.<sup>35</sup> We suggest that content is only 1 aspect of the comprehensive health misinformation challenge on social media. Therefore, relying on textual content alone is not adequate. There is an emerging need to understand health misinformation from more aspects, including the content, the users who are involved, and the social media environment as an interconnected entity.

One of the key approaches to investigate an infectious disease outbreak is to track the trajectory of the epidemic. In this study, we defined the dissemination of a particular piece of (mis)information as a dynamic process in which the original post (e.g., a tweet) is propagated by retweeting in an information-dissemination network (colloquially referred to as “network” hereafter). Retweeting shows that the user recognizes the importance of the original post and is willing to disseminate the piece of information. Therefore, we focused on retweeting as the information dissemination method.<sup>35</sup>

## ABOUT THE AUTHORS

Lida Safarnejad and Yaorong Ge are with the Department of Software and Information Systems, University of North Carolina at Charlotte. Qian Xu is with the School of Communications, Elon University, Elon, NC. Siddharth Krishnan is with the Department of Computer Science, University of North Carolina at Charlotte. Arunkumar Bagarvathi is with the Department of Computer Sciences, Oklahoma State University, Stillwater. Shi Chen is with the Department of Public Health Sciences and the School of Data Science, University of North Carolina at Charlotte.

Correspondence should be sent to Shi Chen, Department of Public Health Sciences, University of North Carolina at Charlotte, 9021 University City Blvd, Charlotte, NC 28223 (e-mail: [schen56@uncc.edu](mailto:schen56@uncc.edu)). Reprints can be ordered at <http://www.ajph.org> by clicking the “Reprints” link.

This article was accepted June 16, 2020.

<https://doi.org/10.2105/AJPH.2020.305854>

We use the 2016 Zika epidemic as a case study when health agencies had prominent presence on social media to share the latest findings and guidelines. However, uncertainty about this epidemic and overlapping events such as the 2016 US presidential election and the Olympics in Rio de Janeiro, Brazil, opened the door for misinformation. In this study, we first established an operational definition of health-related misinformation and identified the most popular Zika misinformation tweets. We then developed an algorithm to infer and reconstruct information dissemination networks of top Zika misinformation and comparable real-information tweets. Afterward, we applied network analyses to extract network structure metrics for both groups. We investigated how network structures differed quantitatively. This study leverages our understanding on the mechanism of health misinformation dissemination on social media and how it might have outcompeted real information. Eventually, the insights from this project will help with the development of more effective health communication strategies on social media against various misinformation.

## METHODS

We chose the entire year of 2016 (January 1–December 31, 2016) as the sampling period for this study. This time period covered the major milestones in Zika epidemic timeline, including the World Health Organization's (WHO's) initial warning of Zika across the Americas, the official declaration of the public health emergency of international concern, and the end of the public health emergency of international concern. Using Zika as the keyword, we collected a total of 3.7 million English-language tweets in and retweets published in 2016 via the Gnip application programming interface through our university's data science program. This data set was complete, including all English-language Zika discussions on Twitter in 2016. Therefore, our data set provided a more comprehensive and unbiased view of the public discourse of Zika on Twitter in 2016.

### Misinformation Identification and Relevant Information Matching

We ranked all original Zika tweets on the basis of the number of received retweets, from

highest to lowest. Following this descending order, we selected the top 5000 most retweeted tweets as the sample pool. Then we established an operational definition of misinformation such that the information in the tweet was not evidence-based. We used peer-reviewed journal articles and conference proceedings, government and health agencies' (e.g., Centers for Disease Control and Prevention and WHO) reports and statistics, and fact-checking Web sites to evaluate the validity of the tweet. After we identified misinformation in the top 5000 pool, we identified another group of comparable tweets with real information based on their posting time and number of retweets. A detailed description of how this definition was operationalized, assurance of reliability, and examples are provided in Appendix A (available as a supplement to the online version of this article at <http://www.ajph.org>).

We acquired metadata of each original tweet as well as all its retweets, including posting date and time, user names involved, and followees–followers information. We used them to track information dissemination, construct networks, and conduct subsequent analyses.

### Constructing Dynamic Information Dissemination Networks

Because Twitter did not explicitly publicize who–retweeted–whom information, we developed an algorithm based on posting or retweeting time and friends–follower relationship to infer and reconstruct the network of retweeting. A detailed technical report on this algorithm is in Appendix B (available as a supplement to the online version of this article at <http://www.ajph.org>).

We also investigated the temporal dynamics of information dissemination by calculating times to receive 50%, 75%, 90%, 95%, and 100% of all retweets in real and misinformation groups. This demonstrated the temporal variability within and between the 2 groups and gave the networks a temporal aspect as well.

### Computing and Interpreting Important Network Metrics

Once the network was constructed for each tweet, we computed and compared

important network metrics highly relevant to information dissemination both within and between the 2 groups. In this study, we extracted total of 9 metrics: network reach (REA), network influence (NIF), diameter (DIA), density (DEN), modularity (MOD), Wiener index (WEI), structural virality (VIR), top out-degree centrality (OUT) score, and top betweenness centrality (BET) score. These metrics quantified and characterized network structures from different aspects and across network levels. Here we provide a succinct description of these metrics; more detailed technical explanations are in Appendix C (available as a supplement to the online version of this article at <http://www.ajph.org>):

- REA of a network measured the number of unique vertices (i.e., unique user accounts in this network).
- NIF (also known as network size) represented number of all vertices. If each user name retweeted exactly once, then REA should equal NIF. A larger difference between REA and NIF indicated that some user names in the network had retweeted more than once.
- DIA was the shortest distance between the 2 most distant vertices in the network.
- VIR measured average distance between all pairs of vertices in the network.
- DEN measured proportion of potential relationships that actually existed in the network.
- WIE was sum of the shortest paths between all pairs of vertices.
- MOD measured likelihood of dividing a network into potential clusters (i.e., subgroups), within which vertices were highly connected, but loosely connected among subgroups. MOD was a local-level network metric.
- OUT measured how much influence a single vertex had in terms of generating retweets to further spread information. We calculated the entire OUT distribution and presented the largest OUT value of all retweeters in a network.
- BET quantified the importance of the vertex in terms of the connectivity of the network. Larger BET indicated a more critical role in maintaining network stability. We showed the largest BET value.

In summary, these network metrics comprehensively characterized different aspects of networks at multiple scales, from overall global network level (REA, NIF, DIA, VIR, DEN, and WIE) to local cluster level (MOD) and all the way down to individual vertex level (OUT and BET). Although there were other network metrics, these 9 metrics were especially critical for information dissemination from original posting user account to other retweeting vertices in the network.

### Network and Statistical Analyses


We compared network metrics between the 2 groups by using the Kolmogorov–

Smirnov test to identify any significant differences in distributions of these metrics. We performed data retrieving, processing, and network reconstruction in Python version 3.7.3 (Python Software Foundation, Beaverton, OR). We carried out network and statistical analyses in R version 3.5.0 (R Foundation, Vienna, Austria) with additional packages. All input data and codes are freely available upon request.

### RESULTS

We focused on the most popular tweets about Zika in this study. We defined

popularity as number of retweets received of a given tweet. We considered a tweet to be popular if it was retweeted at least 50 times. About 5000 tweets in our data set had retweets above this cut off. Among the top 5000 most retweeted Zika tweets in 2016, we identified and verified a total of 400 tweets that contained misinformation. Among them, 266 tweets included adequate metadata to reconstruct the information dissemination networks. Not all metadata were available because of data loss, including user accounts banned by Twitter, content removed by Twitter, or the user actively retracted the original post for various reasons. The comparison group of real information contained a



Note. The y-axis (time) is in a natural logarithm scale. The x-axis (percentage of retweets received) is not positioned by their absolute numbers. On average, misinformation (red) needed a much shorter time to receive 50% of all its retweets compared with real information (black), but a significantly longer time to achieve 90% and higher.

FIGURE 1—Temporal Heterogeneity in Retweeting Activity in Zika Real and Misinformation Groups: 2016

total of 458 tweets that occurred within similar dates of posting of misinformation tweets and had similar number of retweets of misinformation. To avoid potential selection bias, we did not make a 1-to-1 match of real Zika tweets.

### Temporal Variability in Information Dissemination Dynamics

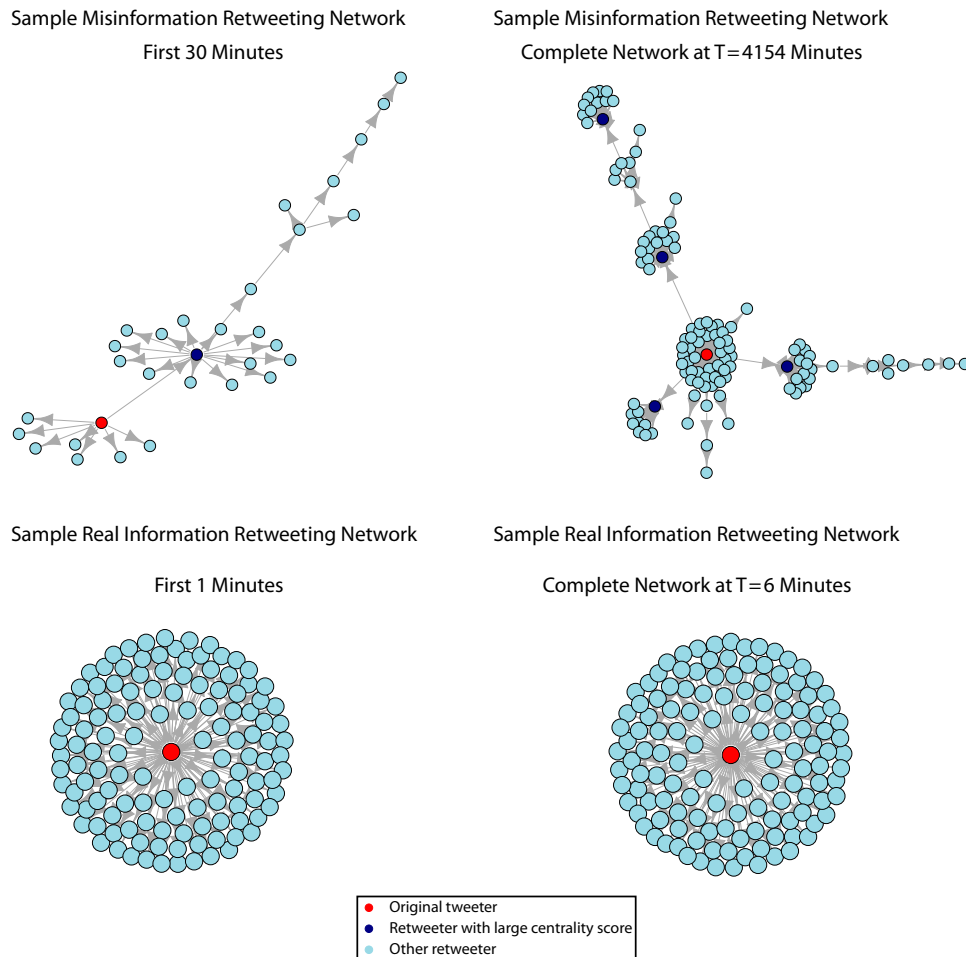
There was substantial temporal variability in the retweeting dynamics between misinformation and real-information groups (Figure 1). It took a significantly shorter time for misinformation to receive 50% of all retweets ( $T_{50} = 334$  min for misinformation;  $T_{50} = 448$  min for real information;  $P < .001$  according to the 2-sided  $t$  test). The difference

was minimal to receive 75% of all retweets ( $T_{75} = 916$  min for misinformation;  $T_{75} = 898$  min for real information;  $P = .93$ ). Interestingly, it then always took a significantly longer time for misinformation to receive 90% of retweets ( $T_{90} = 2580$  min vs  $T_{90} = 1795$  min;  $P = .03$ ), 95% of retweets ( $T_{95} = 4739$  min vs  $T_{95} = 2824$  min;  $P = .001$ ), and all retweets ( $T_{100} = 34\,869$  min vs  $T_{100} = 22\,340$  min;  $P < .001$ ). Misinformation attracted at least half of all retweets within a relatively short period of time to make it more viral. Afterward, misinformation might be deliberately retweeted to keep their visibility over a longer time span. Based on these observations, we chose the time until the last retweet to construct the network, as it provided the most complete view of retweeting activity.

### Network Metrics of Real vs Misinformation Groups

We reconstructed retweeting networks for each tweet in real and misinformation groups. Examples of dynamic network structures are provided in Figure 2 for both misinformation and real information.

The distributions of important network metrics of both groups are presented in Table 1. For demonstration purpose, we scaled all network metrics between 0 and 1 with feature scaling. Actual numeric summary statistics are shown in Table 1. None of these distributions approximated normal distribution, showing high skewness and kurtosis as well as possible multimodality. This indicated large within-group variability of network structures. All network metrics' distributions differed



*Note.* These networks demonstrated how misinformation and real information infiltrated among retweeters. The relative contribution of users could be quantified by their centrality scores. Note that in this example, real information may have been disseminated by potential bots. Therefore, bots are not necessarily associated with misinformation exclusively.

**FIGURE 2—Sample Information Dissemination Through Retweeting Networks of Zika Misinformation and Real-Information Tweets: 2016**

TABLE 1—Statistics of Network Metrics in Zika Misinformation and Real-Information Groups on Twitter: 2016

Network Metric	DIA	WIE	DEN	SIZ	REA	VIR	MOD	BET	OUT
<b>Misinformation group</b>									
Min	1	1	0	0	1	0.5	-0.5	0	1
25%	8	4.10*10 <sup>4</sup>	3.1*10 <sup>-3</sup>	135	136	3.53	0.47	66	43
50%	15	1.14*10 <sup>5</sup>	5.2*10 <sup>-3</sup>	263	264	6.49	0.68	278	89
75%	25	3.95*10 <sup>5</sup>	8.5*10 <sup>-3</sup>	321	322	10.23	0.79	881	148
Max	259	5.19*10 <sup>6</sup>	2.5*10 <sup>-1</sup>	1 456	1 457	60.82	0.93	35 451	844
Mean	20	8.84*10 <sup>5</sup>	9.0*10 <sup>-3</sup>	264	263	8.01	0.61	1 003	115
SE	5	2.26*10 <sup>4</sup>	2.0*10 <sup>-3</sup>	14	13	0.4	0.01	170	7
95% CI	10, 30	4.42*10 <sup>5</sup> , 1.33*10 <sup>6</sup>	5.1*10 <sup>-3</sup> , 1.3*10 <sup>-2</sup>	237, 291	238, 288	7.23, 8.79	0.59, 0.63	670, 1336	101, 448
<b>Real-information group</b>									
Min	1	1	0	0	1	0.5	-0.5	0	1
25%	3	1.54*10 <sup>4</sup>	5.5*10 <sup>-3</sup>	111	112	2.19	0.18	4	57
50%	5	2.71*10 <sup>4</sup>	6.9*10 <sup>-3</sup>	141	145	2.62	0.37	15	97
75%	8	6.02*10 <sup>4</sup>	8.8*10 <sup>-3</sup>	178	179	4.09	0.6	110	131
Max	92	1.31*10 <sup>5</sup>	2.5*10 <sup>-1</sup>	341	342	32.83	0.88	3 567	250
Mean	8	5.80*10 <sup>4</sup>	1.0*10 <sup>-2</sup>	145	144	3.82	0.39	127	97
SE	0.4	4.83*10 <sup>3</sup>	1.0*10 <sup>-3</sup>	3	3	0.15	0.01	15	2
95% CI	7, 9	4.85*10 <sup>4</sup> , 6.74*10 <sup>4</sup>	8.0*10 <sup>-3</sup> , 1.2*10 <sup>-2</sup>	139, 151	138, 150	3.53, 4.11	0.37, 0.41	98, 156	93, 101

Note. BET = top betweenness centrality score of a vertex in the network; CI = confidence interval; DEN = density of follower–followee network (not information dissemination network; see Appendix C, available as a supplement to the online version of this article at <http://www.ajph.org>, for detailed explanation); DIA = diameter; MOD = modularity; OUT = top out-degree centrality score of a vertex in the network; REA = reach; SIZ = size; VIR = virality; WIE = Wiener index. The 2-sided Kolmogorov–Smirnov test was used to investigate whether the distribution of specific network metrics between misinformation and real-information groups were significantly different (i.e., the metrics were from different distributions). All 9 metrics showed statistically significant differences ( $P < .001$ ) between real and misinformation groups.

significantly ( $P < .05$ ) between real and misinformation groups, according to the Kolmogorov–Smirnov test. Therefore, real and misinformation networks presented a lot of heterogeneities, both within and between groups.

The DEN was significantly higher in the misinformation group as retweeters were more likely to engage in retweeting misinformation if their friends (whom they followed) tweeted or retweeted so. Note that this network was the original follower–followee network from which information dissemination network was inferred. The difference and relationship between these 2 networks are detailed in Appendix B. The reason that we used the original network only for density is discussed in Appendix C.

The DIA was also significantly higher in the misinformation group. In general, the smaller the diameter, the fewer layers the information passed through the network to reach the outermost retweeter. Misinformation attracted more grass-roots users one after another, as opposed to more hierarchical, cascade-like dissemination in the real-information group.

The VIR, which focused on average path length, was also significantly higher in the misinformation group, indicating that vertices were generally farther apart in the network. This finding reinforced our previous finding that misinformation involved more direct user-to-user, or small cluster-to-cluster information dissemination than hierarchical dissemination through layers in the real-information group.

The NIF and REA were similar metrics in which REA focused on unique retweeters. The misinformation group had both significantly smaller REA and NIF. In addition, about 30% of misinformation tweets had a same vertex retweeted at least twice, which was substantially higher than in the real-information group (<10%). This could be an intentional propagation strategy to disseminate (mis)information on social media. However, the risk of such strategy was that Twitter might detect it and take actions. Therefore, having multiple user names to retweet the same content together would be a more effective way for information dissemination than having the same user name to retweet the same content multiple times.

For the WIE, the misinformation group had significantly smaller values on average. This indicated that misinformation-retweeting networks could have more star-like local clusters, which reduced the WIE. This was also consistent with our previous finding that the misinformation network had more local clusters and a larger MOD value. However, this finding seemed to contradict the previous finding that, on average, DIA was actually larger in the misinformation group. The actual distribution of WIE in the misinformation group was the key to solve this dilemma (Table 1). For the misinformation group, the WIE distribution had more than 1 prominent peak (i.e., multimodal). While some Zika misinformation networks had overall smaller WIEs, a few networks had very large WIEs. From an actual (mis)information propagation perspective, this implied that propagators exploited 2 seemingly contrasting strategies: first, using a star-like network with a very small WIE (much less frequently observed in the real-information group); second, using chain-like dissemination network with a large WIE. In addition, there were hybrids of these strategies to

disseminate misinformation farther out. For example, propagators of misinformation might create an initial burst of retweets, shown as local stars in the network, which attracted more grass-roots users to help pick up the trend and retweet one after another. However, we did not observe such a sophisticated arrangement in the real-information group.

At the local network level, we saw higher modularity more frequently in the misinformation group, indicating that users who retweeted misinformation tended to form smaller, local clusters to disseminate the information. Therefore, Zika misinformation was more difficult to tackle with traditional mitigation strategies. Multiple smaller clusters reduced the risk of removal of some clusters, as other clusters served as alternative routes for information dissemination. By comparison, the real-information group had relatively smaller MOD. MOD in the misinformation group was more heavily skewed to the left, compared with that in the real-information group (Table 1).

At the individual vertex level, distribution of OUT in the misinformation group also had a strong multimodal pattern, indicating that many misinformation tweets involved a user with an extremely large outbound degree (centrality score >200; Table 1) who might serve as potential online influencer or propagator. On the other hand, distribution of OUT in the real-information group was similar to a normal distribution. OUT of the misinformation group was more heavily right skewed in comparison with the real-information group. The vast majority of users in the misinformation group were not influential in relaying information.

For BET, top BET users in the real-information group had a significantly smaller BET score than that in the misinformation group (127 vs 1003). Therefore, these top BET users in the misinformation group were more important than their counterparts in the real-information group to maintain network stability, as higher BET score indicated a more critical role in information pass-through. While top OUT users could be identified relatively easily by their superficial activity of attracting many retweets, the top BET user, on the other hand, was much more difficult to detect unless constructing the network and performing centrality calculation for all

vertices. Nevertheless, from a misinformation mitigation perspective, targeting top BET users could be a more effective way to stall or even completely shut off misinformation propagation than focusing on top OUT users.

To summarize, the misinformation group had distinct distributions of all of these network metrics from the real-information group, indicating significantly different dissemination network structures. Data mining of information dissemination networks could help health professionals and the general public better understand the dissemination process of health misinformation. In addition, these quantitative metrics could be utilized by health informaticians to develop more accurate infosurveillance and misinformation detection systems.

## DISCUSSION

We developed an analytical framework to investigate health misinformation dissemination on social media. We provided an operational definition of health misinformation and constructed an algorithm to explicitly track how health (mis)information is disseminated on social media through retweeting networks.

We need to point out that our current knowledge about health topics evolves through time as more and more clinical, epidemiological, and other evidence becomes available—hence, the idea of “evidence-based.” The terms “real information” and “misinformation” should be used with caution because our current understanding might be falsified in the future. Timing of the discussion should be considered especially during an emerging health crisis such as the Zika epidemic. For example, we found that *The Economist*, a generally reliable source of information, tweeted in December 2016 that Zika is harmless to adults (the post was deleted) when there had already been clear evidence to show the causal relationship between Zika virus infection and Guillain-Barré syndrome in adults. Had the tweet appeared in early 2016 when the causal relationship was not yet established, it would not have been deemed as misinformation. As a consequence, an important follow-up of this study is to increase the health literacy in the

society such that people learn how to check the validity of health information and why it is misinformation, and frequently update their knowledge about the health issues, instead of merely being told whether a piece of information is real or not.

We will further investigate user activity and attributes to identify bots and examine their effectiveness in dissemination of real and misinformation on social media. The example of a real-information dissemination network (Figure 2) is suspected to be facilitated by bots. In addition, we checked user verification status and only a small fraction were verified users. The real-information group had a higher user-verification rate (2%) than did the misinformation group (1.2%). This observation agrees with other studies showing a strong presence of established news agencies and health organization on Twitter during the Zika outbreak in 2016.

Other work is currently under way to investigate users' activity through time and how this temporal dynamic reveals misinformation infiltration. In this study, we constructed a network  $G$  of a given tweet at the end of all retweeting activities. As we showed large temporal heterogeneity both within and between groups (Figure 1), our algorithm is able to construct a dynamic network  $G_t$  at given time  $t$ . If we detect a sudden rise in retweeting dynamics at time  $t$ , we can then construct a specific network by time  $t$  to explicitly identify which user is causing the burst of retweets, quantify the user's centrality scores, and work at the individual vertex level to further address the health misinformation epidemic on social media.

## Public Health Implications

This study provided solid evidence on health misinformation dissemination patterns on Twitter, one of the most utilized social media platforms. Our analytical framework is universally developed and can explore other public health issues on social media. For example, we have studied genetically modified organism misinformation spreading on Sina Weibo, the largest Chinese social media platform. Other emerging and controversial health topics, including the current COVID-19 pandemic, are also being investigated with this framework.

Another key public health implication of this study is to extract important features of health misinformation, which are not directly identifiable from its content alone. We showed the importance to treat misinformation (pathogen), users (hosts), and social media (environment) as an interconnected entity—the Infodemiology Triad. Misinformation, like real pathogens, is not leaving no trace behind. This study substantially increased our understanding of misinformation dissemination dynamics. Furthermore, the rich data set can be used in conjunction with other features of misinformation (e.g., content, linguistic, and account-based) to build a comprehensive health-misinformation detector. In subsequent work, we will use state-of-the-art machine-learning methods to build such a classifier for public health use.

## Conclusions

We investigated health-related real and misinformation disseminated on social media during a health emergency from a dynamic network perspective. We discovered that the 2 groups had distinct network structures, indicating their different dissemination patterns. Our work has shed light on developing more accurate health misinformation detectors. **AJPH**

## CONTRIBUTORS

Q. Xu, Y. Ge, S. Krishnan, and S. Chen formulated the study idea. A. Bagarvathi collected the data. L. Safarnejad and S. Chen performed analyses and developed the article. Y. Ge and S. Chen supervised the project. All authors contributed substantially in this project.

## ACKNOWLEDGMENTS

This study was supported by the Page Legacy Scholar Grant from the Arthur W. Page Center for Integrity in Public Communication at the Donald P. Bellisario College of Communications at the Pennsylvania State University. This study was also supported by the Faculty Research and Development Sabbatical at Elon University.

## CONFLICTS OF INTEREST

The authors have no conflict of interest in this study.

## HUMAN PARTICIPANT PROTECTION

This study used secondary online social media data in the public domain for analysis. No actual human participants were involved.

## REFERENCES

- Zhao L, Chen F, Dai J, Hua T, Lu C-T, Ramakrishnan N. Unsupervised spatial event detection in targeted domains with applications to civil unrest modeling. *PLoS One*. 2014;9(10):e110206. <https://doi.org/10.1371/journal.pone.0110206>

- Pew Research Center. Social media fact sheet. 2019. Available at: <https://www.pewresearch.org/internet/fact-sheet/social-media>. Accessed January 15, 2020.
- Seymour B, Getman R, Saraf A, Zhang LH, Kalenderian E. When advocacy obscures accuracy online: digital pandemics of public health misinformation through an antifluoride case study. *Am J Public Health*. 2015;105(3):517–523. <https://doi.org/10.2105/AJPH.2014.302437>
- Broniatowski DA, Jamison AM, Qi S, et al. Weaponized health communication: Twitter bots and Russian trolls amplify the vaccine debate. *Am J Public Health*. 2018;108(10):1378–1384. <https://doi.org/10.2105/AJPH.2018.304567>
- Dhaliwal D, Mannion CA. *Anti-vaccine Messages on Facebook: A Preliminary Audit*. Undergraduate Research Report. University of Calgary; 2019.
- Jamison AM, Broniatowski DA, Dredze M, Wood-Doughty Z, Khan D, Quinn SC. Vaccine-related advertising in the Facebook ad archive. *Vaccine*. 2020; 38(3):512–520. <https://doi.org/10.1016/j.vaccine.2019.10.066>
- Sharma M, Yadav K, Yadav N, Ferdinand KC. Zika virus pandemic—analysis of Facebook as a social media health information platform. *Am J Infect Control*. 2017; 45(3):301–302. <https://doi.org/10.1016/j.ajic.2016.08.022>
- Albarracín D, Romer D, Jones C, Hall Jamieson K, Jamieson P. Misleading claims about tobacco products in YouTube videos: experimental effects of misinformation on unhealthy attitudes. *J Med Internet Res*. 2018;20(6):e229. <https://doi.org/10.2196/jmir.9959>
- Chou WY, Prestin A, Lyons C, Wen KY. Web 2.0 for health promotion: reviewing the current evidence. *Am J Public Health*. 2013;103(1):e9–e18. <https://doi.org/10.2105/AJPH.2012.301071>
- Lowry M, Fouse D. Communicating research in an era of misinformation. *Am J Public Health*. 2019; 109(5):645. <https://doi.org/10.2105/AJPH.2019.305048>
- Signorini A, Segre AM, Polgreen PM. The use of Twitter to track levels of disease activity and public concern in the US during the influenza A H1N1 pandemic. *PLoS One*. 2011;6(5):e19467. <https://doi.org/10.1371/journal.pone.0019467>
- Bail CA. Emotional feedback and the viral spread of social media messages about autism spectrum disorders. *Am J Public Health*. 2016;106(7):1173–1180. <https://doi.org/10.2105/AJPH.2016.303181>
- Colditz JB, Chu K-H, Emery SL, et al. Toward real-time infoveillance of Twitter health messages. *Am J Public Health*. 2018;108(8):1009–1014. <https://doi.org/10.2105/AJPH.2018.304497>
- Safarnejad L, Ge Y, Xu Q, Chen S. EventPeriscope: a novel pipeline to analyze impact of real-world events on Twitter discussions. Presented at: International Conference on Social Computing, Behavior-Cultural Modeling and Prediction and Behavior Representation in Modeling and Simulation; July 9–12, 2019; Washington, DC.
- Vraga EK, Bode L. I do not believe you: how providing a source corrects health misperceptions across social media platforms. *Inf Commun Soc*. 2018;21(10): 1337–1353. <https://doi.org/10.1080/1369118X.2017.1313883>
- Gesser-Edelsburg A, Diamant A, Hijazi R, Mesch GS. Correcting misinformation by health organizations during measles outbreaks: a controlled experiment. *PLoS One*. 2018;13(12):e0209505. <https://doi.org/10.1371/journal.pone.0209505>
- Gough A, Hunter RF, Ajao O, et al. Tweet for behavior change: using social media for the dissemination of public health messages. *JMIR Public Health Surveill*. 2017;3(1):e14. <https://doi.org/10.2196/publichealth.6313>
- Park A, Bowling J, Shaw G Jr, Li C, Chen S. Adopting social media for improving health: opportunities and challenges. *N C Med J*. 2019;80(4):240–243. <https://doi.org/10.18043/ncm.80.4.240>
- Bode L, Vraga EK. See something, say something: correction of global health misinformation on social media. *Health Commun*. 2018;33(9):1131–1140. <https://doi.org/10.1080/10410236.2017.1331312>
- Pramanik S, Wang Q, Danisch M, Guillaume J-L, Mitra B. Modeling cascade formation in Twitter amidst mentions and retweets. *Soc Netw Anal Min*. 2017;7(1):41. <https://doi.org/10.1007/s13278-017-0462-1>
- Liang H, Fung IC-H, Tse ZTH, et al. How did Ebola information spread on Twitter: broadcasting or viral spreading? *BMC Public Health*. 2019;19(1):438. <https://doi.org/10.1186/s12889-019-6747-8>
- Zhang Q, Zhang S, Dong J, Xiong J, Cheng X. Automatic detection of rumor on social network. In: Li J, Ji H, Zhao D, Feng Y, eds. *Natural Language Processing and Chinese Computing*. Cham, Switzerland: Springer; 2015: 113–122. [https://doi.org/10.1007/978-3-319-25207-0\\_10](https://doi.org/10.1007/978-3-319-25207-0_10)
- Hamidian S, Diab M. Rumor identification and belief investigation on Twitter. In: *Proceedings of the 7th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*. San Diego, CA: Association for Computational Linguistics; 2016. <https://doi.org/10.18653/v1/W16-0403>
- Thakur HK, Gupta A, Bhardwaj A, Verma D. Rumor detection on Twitter using a supervised machine learning framework. *Int J Inf Retr Res*. 2018;8(3):1–13. <https://doi.org/10.4018/IJIRR.2018070101>
- Jin F, Wang W, Zhao L, et al. Misinformation propagation in the age of Twitter. *IEEE Computer*. 2014;47(12):90–94. <https://doi.org/10.1109/MC.2014.361>
- Soe SO. Algorithmic detection of misinformation and disinformation: Gricean perspectives. *J Doc*. 2018; 74(2):309–332. <https://doi.org/10.1108/JD-05-2017-0075>
- Yu F, Liu Q, Wu S, Wang L, Tan T. A convolutional approach for misinformation identification. In: *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence*. International Joint Conferences on Artificial Intelligence; 2017. <https://doi.org/10.24963/ijcai.2017/545>
- Ghenai A, Mejova Y. Catching Zika fever: application of crowdsourcing and machine learning for tracking health misinformation on Twitter. Presented at: 5th IEEE International Conference on Healthcare Informatics; August 23–26, 2017; Park City, UT: 518.
- Castillo C, Mendoza M, Poblete B. Information credibility on Twitter. In: *Proceedings of the 20th International Conference on World Wide Web*. New York, NY: Association for Computing

- Machinery; 2011. <https://doi.org/10.1145/1963405.1963500>
30. Jain S, Sharma V, Kaushal R. Towards automated real-time detection of misinformation on Twitter. Paper presented at: 2016 International Conference on Advances in Computing, Communications and Informatics (ICACCI); September 21–24, 2016; Jaipur, India. <https://doi.org/10.1109/ICACCI.2016.7732347>
31. Qi S, AlKulaib L, Broniatowski DA. Detecting and characterizing bot-like behavior on Twitter. Paper presented at: International Conference on Social Computing, Behavioral–Cultural Modeling and Prediction and Behavior Representation in Modeling and Simulation; July 10–13, 2018; Washington, DC. [https://doi.org/10.1007/978-3-319-93372-6\\_26](https://doi.org/10.1007/978-3-319-93372-6_26)
32. Wang Y, Ma F, Jin Z, et al. Eann: Event adversarial neural networks for multi-modal fake news detection. In: *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. London, UK: Association for Computing Machinery; 2018.
33. Chadwick A, Vaccari C. News sharing on UK social media: misinformation, disinformation, and correction. Loughborough, UK: Loughborough University; 2019.
34. Chou WS, Oh A, Klein WMP. Addressing health-related misinformation on social media. *JAMA*. 2018; 320(23):2417–2418. <https://doi.org/10.1001/jama.2018.16865>
35. Wu L, Liu H. Tracing fake-news footprints: characterizing social media messages by how they propagate. In: *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*. New York, NY: Association for Computing Machinery; 2018. <https://doi.org/10.1145/3159652.3159677>