Chapter 3

Diffusion of Misinformation: Topological Characteristics and User Vulnerability

Sibo Wang

https://orcid.org/0000-0003-1892-6971

Hai Liang

https://orcid.org/0000-0003-1779-9552

Abstract

This chapter takes Weibo as a case study to illustrate the dynamics of infodemic diffusion during the COVID-19 pandemic with a focus on the topological characteristics of the diffusion pattern, the profile information of participating users and the node types they represent, and the textual features of the information content. Using computational methods, we further analyze the relationship between social media user characteristics (e.g., location, verification status, the number of followers) and their varying levels of vulnerability to misinformation messages related to COVID-19. The findings show that the emergence of COVID-19 misinformation on Weibo is closely related to the severity of the epidemic in China. The dominant sources of misinformation tend to be unverified users. The spread range of the misinformation is smaller than that of true news, but the spread is deeper and has been persistent since the initial outbreaks. Most of the super-spreaders during the propagation process are the source post publishers; the most common grassroots users seem to be the most vulnerable group to COVID-19 misinformation.

Keywords: Diffusion pattern, propagation structure, user characteristics, modeling, cascade

Introduction

With the outbreaks of the COVID-19 pandemic, people's daily life has been restricted by a variety of anti-epidemic measures, such as lockdown, quarantine, and work-from-home, which has led to more screen time and more social networking. People also read, retweet, and discuss COVID-19-related news more often than any other topics, such as questions about all aspects of this epidemic, where it came from, how it transmits, what symptoms it has, where it’s happening, and how should we respond among other things. An ocean of relevant information has emerged. Not only do public health authorities release on social media the latest policies and updates on the latest epidemic situation, but social media platforms also make it possible for ordinary users to post or publish content they create to discuss and spread related news and updates. As a result, a tsunami of COVID-19-themed “infodemic” has swept online social networks.

Since the outbreak of COVID-19 in late 2019, news about the disease has tended to receive very high attention, and thus it is easy to be hot spots. Discussions of COVID-19 topics have become a gathering place for various kinds of intentionally or unintentionally fabricated misinformation about the disease. Some COVID-19 misinformation exaggerates the severity of the epidemic and creates a surge of anxiety and worries in the general public. Refer to Chapters 1 and 2 for specific examples of such misleading and inaccurate messages. As illustrated in Chapter 1, misinformation has brought different degrees of harm to people's well-being and health.

Stimulated by the large user base of social media and the convenience of sharing, resistance to the dissemination of this misinformation is unprecedentedly small. Exposed to a large amount of unsubstantiated information and the opinions of others, users are often affected and become contributors to the further diffusion of the information (Wang & Zhang, 2022; Zhou & Zafarani, 2019). Compared with the rigorous verification of the veracity of the information itself, the influence of the source publisher or disseminators of the information on social media is more likely to have an impact on the judgment of users. A piece of widely disseminated misinformation may be released by a for-profit marketing account with a large fan base. Ordinary users can help spark it by making the news sound convincing and forward to others. In other cases, users who get no attention can use social bots to carry out the early diffusion and increase the credibility of the fake story.

The reason why a piece of true news from an ordinary user is widely disseminated may be that the source news is retweeted by an Internet influencer (i.e., Key Opinion Leader—aka KOL), which gives it more exposure, or it may also be that it tells a truth that readers are eager to know. This sort of diffusion path exactly reflects a unique feature of decentralized social media that is different from traditional news websites with a top-down propagation structure. In the above example, the profile information of users (the number of followers, verification status, the number of historical tweets, etc.) and different degrees of the scope of message transmission belong to propagation structure information. As the explicit embodiment of the different degrees of attractiveness finally brought about by the content and writing styles, it can complement the traditional misinformation detection based on textual feature modeling and become a powerful pillar for identifying misinformation.

Building on the previous chapter concerning the source and message characteristics of COVID-19 misinformation on Weibo, this chapter explores the distinguishable characteristics of misinformation in the propagation structure that are different from that of true information. To illustrate, suppose you are surfing social media and come across a widely retweeted message with a picture that says: "The Russian government has put 800 lions and tigers on the street to prevent people from gathering during the pandemic." Some questions will slide into your mind. It sounds ridiculous but is it true? Who posted this news? How do people react when they retweet? With a skeptical attitude, you will check the veracity of this post carefully and find the truth that Russia did not release any tigers and lions to combat the pandemic. The post has been accompanied by a picture of a film set in Johannesburg, South Africa, in 2016.

This example raises a set of new questions. How can such a piece of misinformation be reposted? Who is the super-spreader? What types of users are vulnerable to misinformation? Can researchers judge misinformation by modeling its existing propagation structure? Chapter 3 will address these major questions. In addition, using Weibo as a case study for the same reasons as Chapter 2 (e.g., lack of access to big data from Facebook or Twitter in Hong Kong, Singapore, and Taipei), we are interested in exploring the differences in these aspects between true and false information about COVID-19, as well as the differences in the topological characteristics of information dissemination on other topics during the same period compared with the epidemic topic.

In short, by constructing a dataset of true and false messages on Weibo during the pandemic, we pursue a comprehensive analysis of misinformation diffusion (intentionally or unintentionally, see Wang & Zhang, 2022), with a focus on the topological characteristics of the diffusion pattern, the profile information of participating users, and the textual features of the information content. Additionally, a machine learning classifier that can accurately identify the veracity of information about COVID-19 and generate the crucial propagation substructures is obtained by modeling the above aspects of features.

Modeling the Topological Structure of Misinformation: An Example

Problem definition

We start with an equation to build the topological structure. $G = (V, E)$ is a directed graph representing an online social network, where V is the node set, including users on the social network, each node represents a social media user, and E is the edge set, including information transmission behavior on the social network, and each edge represents a repost or a comment. We use the cascade $c = (V_c, E_c)$ as the input of the model, which is a connected directed graph representing the whole propagation process of a piece of news n_c, where $Vc \subseteq V$ represents the set of all participants in the propagation process of news n_c in the social network G, f_u^{text} represents the textual feature vector of the post published by user u, f_u^p denotes the profile information features of user u, and $Ec \subseteq E$ denotes the set of edges passed during the propagation of news n_c.

We now illustrate this definition with an example: A piece of misinformation about Dr. Zhong Nanshan, a well-known Chinese anti-epidemic expert, published by a big V user (e.g., KOL) on Weibo with more than 650,000 followers on June 18, 2020. It claimed that Dr. Zhong flew to Beijing to help fight a new wave of epidemic outbreaks in Beijing. The post triggered 599 reposts and comments. A portion of the propagation behaviors in this cascade is shown in Figure 3.1.

[Insert Figure 3.1 Here]

The truth: This message used Zhong Nanshan's expired video, falsely claiming that he had parachuted into Beijing, although he was in Guangzhou at that time. Consequently, the blogger was deducted 10 credit scores for publishing misinformation and was banned for blogging for 15 days. In Figure 3.1, we draw each node in the cascade as a block to facilitate representing the characteristics of the node. The left side of each block lists the profile information of the author of the microblog, that is, f_u^p, including the number of his/her followers, the number of his/her followees (if user *u* follows user *v*, then *u* is the follower of *v* and *v* is the followee of *u*), his/her verification status, his/her verification reason, etc. On the right side is the original microblog.

To analyze the original text, we need to use natural language processing technology, by dividing sentences into words and performing the frequency statistics on the vocabulary, so as to constitute the textual feature of the microblog, namely f_u^{text}. The arrow in Figure 3.1 represents the edge we defined, that is, the flow of information, and the arrow pointing from node *u* to node *v* indicates that user *v* reposted or commented on the user *u*'s post.

Topological structure and cascade effect

By observing the substructure of the cascade drawn in Figure 3.1, we can get some statistics. For example, the maximum depth of the structure is 4, that is, the maximum length of the diffusion chain is 4, and the maximum breadth is 8, that is, a post is directly propagated at most 8 times. All expressed in terms of concepts in graph theory, the node with the maximum outdegree has an outdegree of 8.

Similarly, we analyzed the complete cascade of the misinformation and visualized it in Figure 3.2, where each point represents the node of a microblog, and each arrow represents a directed information dissemination edge. The larger the area of a point, the greater the outdegree of the node. Users are also classified according to the user profile information represented by each node, and the detailed criteria will be explained later. The longest propagation path is also 4, and the node with the largest outdegree is the source post represented by the first block in Figure 3.1. The maximum outdegree is 568. That is, this microblog has been reposted and commented for 568 times, which is also the reason for the radial shape of the entire graph. Other than the source node, there are 20 nodes with outdegree greater than 0, indicating that 20 retweets generate information dissemination later. Using the two Chinese characters “谣” and “假” (rumor and fake) as keywords, we found that in the original cascade, 27 nodes questioned the authenticity of the source post like 4 microblogs at the bottom of Figure 3.1.

[Insert Figure 3.2 Here]

The above example illustrates that the introduction of topological structure in studying misinformation not only provides a perspective other than content to model it, but also reveals more supplementary information in terms of text. We can capture evidence about the veracity of posts from different opinions about it expressed by all participants during its propagation. For the sake of comparison, we use true and false as the veracity expression of our two categories of information online, because fake news is more like a political metaphor.

Graph Structured Data and Analyses

Dataset

To study the characteristics of the topological structure of the misinformation related to COVID-19, we used Sina Weibo’s API to collect and build a new dataset with a period from November 2019 to March 2022, including 4,174 source posts, which added up to 961,962 microblogs together with reposts from these sources. The user profile of the author of each microblog was also collected, including the number of followers, number of followees, verification status, verification type, verification reason, description, gender, location, etc. In addition, we scraped the number of historical posts by the users represented by root nodes. The misinformation came from Weibo Community Management Center (a service where users can report a microblog that contains false information), and social media posts collected according to COVID-19 misinformation on social networks published by China Internet Joint Rumor Suppression Platform. The true messages came from verifiably accurate items on official government accounts, the rumor-refuting microblogs provided by Weibo Community Management Center and microblogs posted by other users whose content and release time were consistent with the corresponding messages in the previous two sources.

Since we also sought to explore whether the dissemination characteristics of COVID-19 misinformation differed from those of other topics during the pandemic, this dataset included microblogs of other topics and their propagation cascades in the same period. We not only cared whether there was a difference in the diffusion pattern of the true and false news of COVID-19 that could help us distinguish their veracity, but also whether there was a difference between the spread of misinformation on COVID-19 and the spread of misinformation on other topics. Therefore, our dataset contained 2,171 cascades about COVID-19 and 2,003 cascades about other topics. In the same way, the misinformation came from Weibo Community Management Center, and the true information consisted of the corresponding rumor-refuting microblogs provided by Weibo Community Management Center and microblogs sampled from true messages collected according to Internet hot words during the pandemic. The details of the dataset are shown in Table 3.1, which is the largest Weibo dataset in the pandemic with comprehensive propagation structures of information.

[Insert Table 3.1 Here]

Analyses of Graph Structured Data

According to the above statistical results in Table 3.1, the spread scope of true information was larger than that of false information, especially in the context of COVID-19. The size of cascades of true information was even five times that of false information, which leads us to regard topological structure as an important basis for judging misinformation. In addition, we also visualized the relationship between the amount of misinformation and the time span in the dataset by month. As Figure 3.3 shows, when COVID-19 was first known to the general public in January 2020, a sense of panic erupted on social media. Accompanied by the severity of the epidemic and the spread of cases, the amount of misinformation surged to a peak in March. This is consistent with the pattern in Chapter 2.

As the subsequent epidemic situation gradually came under control, misinformation rarely reappeared, although occasionally there was a small amount coinciding with a small rebound of the epidemic. Due to outbreaks of the highly contagious Omicron variant in China, the pandemic deteriorated sharply at the end of 2021, resulting in an amount of COVID-19 misinformation that reached a peak in January 2022. In other words, the emergence of COVID-19 misinformation on Weibo was closely related to the severity of the epidemic in China; when the epidemic was severe, the government generally adopted a strict lockdown policy, and it was difficult for the general public to rebut the misinformation on the spot, thus allowing rumormongers to send out misinformation. As Zhou & Zafarani (2019) pointed out, rumormongers tend to choose the time when people are most panicked, which further intensifies people's anxiety, makes them sensitive and weakens their ability to discriminate the authenticity of the information, thus helping to spread misinformation.

[Insert Figure 3.3 Here]

Analyses of Diffusion Patterns

The characteristics of cascades of misinformation on Weibo in various aspects of its structural characteristics are shown in Table 3.2. Consistent with Table 3.1, in terms of the cascade size, the spread range of the misinformation was smaller than that of true news, but the dissemination of COVID-19 misinformation had been persistent following the initial outbreaks. For the longest path owned by a cascade, this metric also indicated the maximum depth of each cascade. The longest path of misinformation was longer than that of true information, suggesting that the spread of misinformation went deeper. That is, spreaders of misinformation were more likely to retweet comments of the source message than spreaders of true messages. The smallest standard deviation indicated that cascades of misinformation were the most stable in terms of maximum depth. In cascades, the nodes with the largest degree were generally the root nodes, and non-root nodes accounted for less than 4% of the four types of information.

Row "%Not root node" of attribute "Max Degree" in Table 3.2 represents the proportion of the most retweeted nodes that were not the root node in each category of information; it indicates that most of the super spreaders during the propagation process were the source post publishers, and only a small percentage of information spread more widely because it was reposted by others. On this metric, different from other topics, the probability that the node with the maximum degree of COVID-19 misinformation was not the root node was less than the true information. Statistics on the number of nodes with a degree greater than 1 in the cascade and their proportion in all nodes can reveal how many nodes further diffuse information. The proportion of nodes with further diffusion in misinformation was higher than that of true information.

[Insert Table 3.2 Here]

Who is the Source?

In most cases, the authors of the microblogs represented by the root nodes of propagation cascades were the biggest spreaders in the entire propagation process. Specifically, we visualized the number of followers, followees, and historical posts of users of each cascade root node. For the sake of easy observation, we calculated the natural logarithm of these data and scaled the data to a certain range. As shown in Figure 3.4(a), the horizontal axis represents the number of followers, the vertical axis represents the number of followees, the area of each point represents the number of historical posts, and different markers of points represent different user verification statuses.

[Insert Figure 3.4 Here]

Message publishers of COVID-19 misinformation have an average of 340,989 followers, 1,314 followees and 19,019 historical microblogs, while publishers of COVID-19 normal messages have an average of 16,941,965 followers, 1,109 followees and 56,022 historical posts. According to the number of followers, it is clear that unverified users were mostly distributed in the left half of the figure, while verified users normally had larger orders of magnitude of followers. In contrast, misinformation has more unverified users in the right half of the figure (misinformation messages have more unverified publishers who have large fan bases than true messages), true information has more verified users in the left half of the graph (true messages own more verified publishers who have few followers than misinformation messages), and most of the publishers of true information are concentrated in the right half (most of the publishers of true messages have large fan base). With the increase in the order of magnitude of followers, the number of historical posts owned by users shows an overall upward trend, among which the average number of historical posts published by misinformation publishers is less than that of true information.

Do the source users during the propagation of COVID-19 misinformation have distinguishable geographical distribution characteristics? We carried out frequency statistics for the source users. On a scale from 1 to 100, Beijing (21.25%) was at the top of all provinces, followed by Guangdong (10.63%), Other (8.72%), Shanghai (8.17%), and Overseas (7.36%).

Further, we classified user nodes according to the number of followers, the number of followees, verification status, and account status. In line with the focus on diffusion patterns in this chapter, we took the dissemination potential reflected in each user's profile as a basis for classification, rather than the account type which has been analyzed in Chapter 2. To be specific, the sources with more followers can reach more users, while a microblog of a user with few followers had very low exposure. Consequently, we used the number of followers as one measure of transmission potential. The number of followees shows how widely a user receives posts from others. For example, a user who follows 2,000 accounts is more likely to see a particular message than someone who follows 100 accounts. Combining the two features, an account with a large number of followers but a much lower number of followees generally plays the role of an opinion leader on social media, such as *People's Daily* which has 150 million followers but only follows 3,000 accounts. Therefore, the ratio of the number of followers to the number of followees is also one of our dividing criteria. Besides, personal information is an important means to regulate users' behavior on the Internet for fear of being held accountable for inappropriate comments. The verification status, which requires the submission of real personal information to apply for, makes authors cautious in expressing their opinions and makes readers pay different attention to the judgment of microblogs posted by authors with different verification statuses. Similar to the processing method in Figure 3.4(a), we show the relationship between source user profile information and diffusion cascade size in Figure 3.4(b). It is clear that the area of the points near the upper right is larger, that is, messages posted by users with stronger influence potential are more widely spread. In addition, the points with a large area are generally verified users, indicating that messages disseminated widely, whether true or false, are more likely to be published by verified users.

In summary, the node type of the currently banned account is 0. For an unverified user, if the number of followers is greater than 500 and the number of followers is more than twice the number of followees, then node type 1, which represents users who tend to output opinions and have the potential to be influencers, namely unverified influencers, otherwise it is node type 2, which represents the grassroots users who tend to receive opinions; for a verified user, if the number of followers is greater than 10,000 and the number of followers is more than 20 times the number of followees, then the classification is node type 3, which represents the key opinion leaders (KOLs) such as government accounts and Internet celebrities, otherwise it is type 4, which represents the verified individual account. This classification method helps us to preliminarily divide users' influence potential on social media from the aspect of their profile information. Based on these user profile characteristics, we claim that type 3 users (KOLs) have the strongest influence potential, the weakest is type 2 (grassroots users), and type 0 (banned users) is not involved in the comparison as a specific category due to no accessible profile information. We visualized the distributions of user types under various conditions in Figure 3.5. Note that because some nodes included in a cascade may be from the same user, we used the user type of the author of each node to analyze the participation degree of this type during propagation, instead of only considering unique users. For instance, a cascade of misinformation contains 10 tweets, but eight of them are from two grassroots users and only two are from two KOLs. Then the 1:1 result of these two user types can be calculated by the number of unique users, which misleads our judgment on the engagement of different types of users.

We then compared the node type distribution of publishers. As Figure 3.5(a) shows, except for COVID-19 misinformation, misinformation of other topics was mainly published by the most influential KOLs, while the users with the largest proportion of the publishers of COVID-19 misinformation were the least influential grassroots users. More than half of other kinds of information was posted by verified users, while COVID-19 misinformation was significantly smaller in the proportion of verified users. It is worth noting that when checking all the source user accounts in July, we found that 134 were already banned.

It is difficult to generalize about the main content of unverified users, but the verification reasons provided in verified accounts that make up the majority of sources can help us further understand COVID-19 misinformation rumormongers. We conducted keyword statistics on their verification reasons and found that We-Media was the majority, and their usual microblogs were related to star chasing, video, reading and writing, finance, science and technology, entertainment, etc. Except for medicine bloggers, there was little difference in the ability of bloggers in other fields to identify misinformation.

Who Is the Most Vulnerable?

Table 3.3 shows the proportions of node types of all tweets and unique users in the whole dataset, from which it is clear that grassroots dominated (79.23%) the category while KOL accounted for the least except for the banned users. This not only reveals that our dataset reasonably sampled the Weibo platform—where KOLs are a minority and grassroots users are the most common—but also reflects the different contributions of different types of users to the formation of the propagation cascades, which will be elaborated in combination with Figure 3.5(b).

[Insert Table 3.3 Here]

For geographical distribution of all participants in COVID-19 misinformation propagation, Other (21.03%), Beijing (9.96%), Guangdong (9.74%), Shanghai (6.17%), and Overseas (5.53%) ranked in the top five regions on a scale from 1 to 100. The pattern shows the tendency of disseminators who hid their real identifities.

According to Figure 3.5(b), grassroots users are indeed the most common type during the dissemination of information on social media, which occupies a dominant position regardless of the spread of true or false information. However, the proportion of grassroots users in the dissemination of misinformation is also significantly higher than that in true information. It can be concluded that the most common users on social media are the most vulnerable group and the main component of the dissemination structure of COVID-19 misinformation. The whopping 82% percentage of type 2 reveals that an overwhelming number of grassroots users participate in the cascade of misinformation dissemination. As the receivers of information, they can easily repost or comment on the microblogs they see, and the authenticity of the messages does not affect their information transmission behavior.

KOLs with the greatest influence represent the smallest proportion in the dissemination of each kind of information. As opinion leaders, they rarely participate in the further propagation of information but usually release the source microblogs as shown in Figure 3-5(a, b). Occupying second place in the misinformation dissemination user types are unverified influencers with greater influence among unverified users. Second place in true information dissemination goes to the less influential set of individual verified users. The proportion of verified users in the dissemination of misinformation is significantly smaller than that in the dissemination of true information. We can see how much influence can reflect how cautious users are about spreading behaviors. Verified users, as users with real-name information endorsements, generally have higher credibility and are more inclined to disseminate true information.

Table 3.2 shows that the root node was also the largest spreader in most cascades. Results obtained by directly counting the largest spreaders were the same as those shown in Figure 3.5(a). Therefore, we observed the node type distribution of both the root node and the largest spreader in the nodes other than the root node. Cascades smaller than 3 were excluded. As shown in Figure 3.5(c), grassroots users made the biggest contribution to the spread of COVID-19 misinformation; their proportion of 39.40% indicates that they were the super-spreaders of COVID misinformation. This is a very interesting phenomenon, because generally the microblogs that can be further spread are posted by users with enough fans or have rich textual characteristics, for example, they fully express personal views or are trolling. However, we note that most of the microblogs published by grassroots users contained only a few simple words. What causes such microblogs to get a lot of dissemination is worth further investigation.

At the same time, Figure 3.5(c) again reflects the opinion leader positioning of type 3 users. Although the proportion of KOLs decreased compared with the node type distribution of the root nodes in Figure 3.5(a), it still dominated the super-spreaders of the latter three kinds of information. Overall, there were more unverified users among the largest spreaders of misinformation, while there were more verified users among spreaders of real information. The largest spreaders in nodes other than the root node had a certain proportion of banned users, especially the highest proportion in COVID-19 misinformation. It should be noted that Weibo occasionally punishes misinformation spreaders by deducting credit scores. Banned users are usually a particular group of users whose statements involve sensitive topics and have a great impact, which are considered by Weibo authorities to have damaged the harmony of the online community. The proportion of type 0 shows the emergence of super spreaders who made comments that Weibo administraters considered very bad for hamony in the process of information transmission. For example, the account of Wang Sicong, one of the most famous rich second generation in China, was suspended after he published remarks that COVID-19 testing was actually a test of servility and questioned the effectivenesss of Lianhua Qingwen Capsules.

[Insert Figure 3.5 Here]

To further explore the characteristics of COVID-19 misinformation diffusion patterns and user type distribution, we developed a Graph Convolutional Network model to combine and exploit the textual feature and topological structure of cascades in our Weibo-COVID-19 dataset, to predict the diffusion pattern and veracity, and finally achieve a good result of 93% accuracy in detecting misinformation task.

Our model simulated the crucial diffusion substructure by selecting a subset of important nodes in the complete propagation cascade. We visualized the node type distribution of all the participants in generated propagation patterns in Figure 3.5(d). Compared with the complete cascade shown in Figure 3.5(b), grassroots users remained the largest disseminator group, but the proportion decreased, while the proportion of verified users, especially KOLs, showed the largest improvement. When selecting nodes, our model tended to select nodes with more text information under the consideration of node features. As mentioned above, many grassroots users did not express their views on the source information, but just reposted microblogs. However, KOLs generally expressed their opinions on the source information, so the probability of being selected during training increased. The grassroots user base was large enough and consequently there were enough text-informative candidates to choose from, so this only caused minor fluctuations in the proportions of the node types, which further confirmed that our model well captured the importance of different user types in determining cascade substructures whose information belongs to different categories.

Summary of Key Findings

Using big data from Weibo and computational methods and with a focus on identifying the topological characteristics of the diffusion pattern of COVID-19 misinformation as well as profiling Weibo users and the node types they represent, a number of patterns emerged in the present chapter. They shed lights on how users propagate COVID-19 misinformation. Those patterns were compared with the dissemination of true messages. In summary,

- The emergence of COVID-19 misinformation on Weibo was closely related to the severity of the epidemic situations in China. With each outbreak, such as the early stage of the pandemic and the emergence of Omicron, the amount of misinformation soared to a peak. On the other hand, misinformation rarely appeared when the epidemic situation was under control.
- The spread range of COVID-19 misinformation was smaller than that of true messages, but the spread of misinformation went deeper (i.e., more dissemiantors of true messages retweeted the source microblog directly than that of misinformation).
- The sources concerning microblogs of COVID-19 misinformation spreading on Weibo were mostly published by KOLs and a large number of grassroots users; and they were mainly distributed in three of China’s largest and most important cities: Beijing, Guangdong, and Shanghai. The largest proportion of misinformation sources were the least influential users (unverified with few followers).
- Most of the super-spreaders were the source post publishers, and the probability that the largest disseminator of COVID-19 misinformation was not the source publisher was less than that of true messages. The proportion of users who caused further retweets in misinformation was higher than that in true messages.
- Grassroots users were the most vulnerable group to COVID-19 misinformation. Using

real-name and verification may reduce users' reposting unchecked misinformation messages on social media.

Conclusion and Insights

Leveraging the reposting function on Weibo, in this chapter, we tracked and reconstructed the complete diffusion paths of thousands of misinformation and normal messages. By doing so, we are able to identify the unique topological characteristics in the diffusion patterns of misinformation messages.

First, the source microblogs of COVID-19 misinformation spreading on Weibo were mostly published by KOLs and grassroots users and were mainly distributed in Beijing, Guangdong, and Shanghai. Consistent with a previous study (Goel et al., 2016), only a small proportion of messages spread widely and deeply through reposting. Most super-spreaders in the diffusion cascades were the source accounts, mainly including KOLs and grassroots users. This implies that policymakers should focus on the sources to cope with misinformation spreading. However, as we found in chapter 2, it might be difficult to identify misinformation sources in advance because most do not include any external URLs and many of the accounts are government and media accounts. One possible way to solve the problem is to improve the misinformation literacy of local government and media account managers. In addition, it is also consistent with a previous study (Vosoughi et al., 2018) that misinformation messages generally spread deeper than do normal messages.

Second, given the importance of sources, we further analyzed their characteristics. The findings indicate that the largest proportion of misinformation sources are the least influential users (unverified with few followers). Given the sheer volume of these grassroots users, it might be difficult to identify them using automatic approaches.

Finally, our analyses also suggest that grassroots users are the most vulnerable group to COVID-19 misinformation. This is unsurprising because they are the majority of Weibo users. However, it demonstrates that purely focusing on the sources is insufficient to cope with misinformation. It might be equally important to improve the misinformation literacy of the general public. Our findings also indicate that using real-name and verification may reduce users' reposting misinformation messages.

Then, how does the diffused misinformation surrounding COVID-19 spread among the public? How much do people actually view and contribute to the spread by sharing? Also, in what ways do the exposure to and sharing of COVID-19 misinformation differ by social media users' demographics, social differentiators, and the city in which they live? Using large-scale survey data collected in the four studied cities, we examine these empirical questions in the next two chapters.

References

Goel, S., Anderson, A., Hofman, J., & Watts, D. J. (2016). The structural virality of online dif fusion. *Management Science*, *62*(1), 180-196. https://doi.org/10.1287/mnsc.2015.2158

Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science*, *359*(6380), 1146-1151. http://doi.org/10.1126/science.aap9559

Wang R. & Zhang H. (2022). Who spread COVID-19 (mis)information online? Differential informedness, psychological mechanisms, and intervention strategies. *Computers in Human Behavior*, *138*, 107486. https://doi.org/10.1016/j.chb.2022.107486

Zhou, X., & Zafarani, R. (2019). Network-based Fake News Detection: A Pattern-driven Approach. *ACM SIGKDD Explorations Newsletter 21*, 2 (2019), 48–60. https://doi.org/10.48550/arXiv.1906.04210

Dataset **Attribute**	**Weibo**	**Weibo-COVID-19** (**COVID-19** subset)	**Weibo-Other** (Other topics subset)
#cascades	4,174	2,171	2,003
#microblogs	961,962	536,719	425,243
#microblogs per cascade	230	247	212
#cascades of true information	2,087	1,403	684
#cascades of false information	2,087	768	1,319
#microblogs in true cascades	650,600	484,152	166,448
#microblogs per true cascade	312	345	243
#microblogs in false cascades	311,362	52,567	258,795
#microblogs per false cascade	149	68	196

Table 3.1. Details of the dataset built from Weibo

	COVID-19 False	COVID-19 True	Other False	Other True
Cascade Size				
Mean	68.45	345.08	196.21	243.35
Min	1	1	1	1
Median	6	4	11	9
Max	9340	26235	26535	30791
Longest Path				
Mean	2.91	2.76	3.40	3.14
Std	1.95	1.97	2.59	2.41
Min	1	1	1	1
Median	2	2	3	2
Max	15	16	29	17
Max Degree				
Mean	36.66	294.31	83.24	158.90
Min	0	0	0	0
Median	4	3	8	6
Max	2640	22063	7805	22348
%Not root node	3.65	3.99	3.26	2.78
Number of Nodes with Degree>1				
Mean	9.43	14.53	27.30	21.54
Min	0	0	0	0
Median	1	1	2	1
Max	1363	1210	3705	1954
Percentage	13.78%	4.21%	13.91%	8.85%

Table 3.2. Structural characteristics statistics of COVID-19 misinformation on Weibo

Node type	0	1	2	3	4
User type	Banned users	Unverified influencers	Grassroots users	Key opinion leaders	Individual verified users
Tweets in the whole dataset					
Count	11320	70900	762190	32823	84729
Percentage	1.18%	7.37%	79.23%	3.41%	8.81%
Unique Users in Weibo-COVID-19					
Count	4229	27264	330476	11265	36462
Percentage	1.03%	6.66%	80.66%	2.75%	8.90%
Unique Users in Weibo-Other					
Count	6046	24622	274351	6625	24242
Percentage	1.80%	7.33%	81.68%	1.97%	7.22%

Table 3.3. Proportions of node types of all users in the dataset

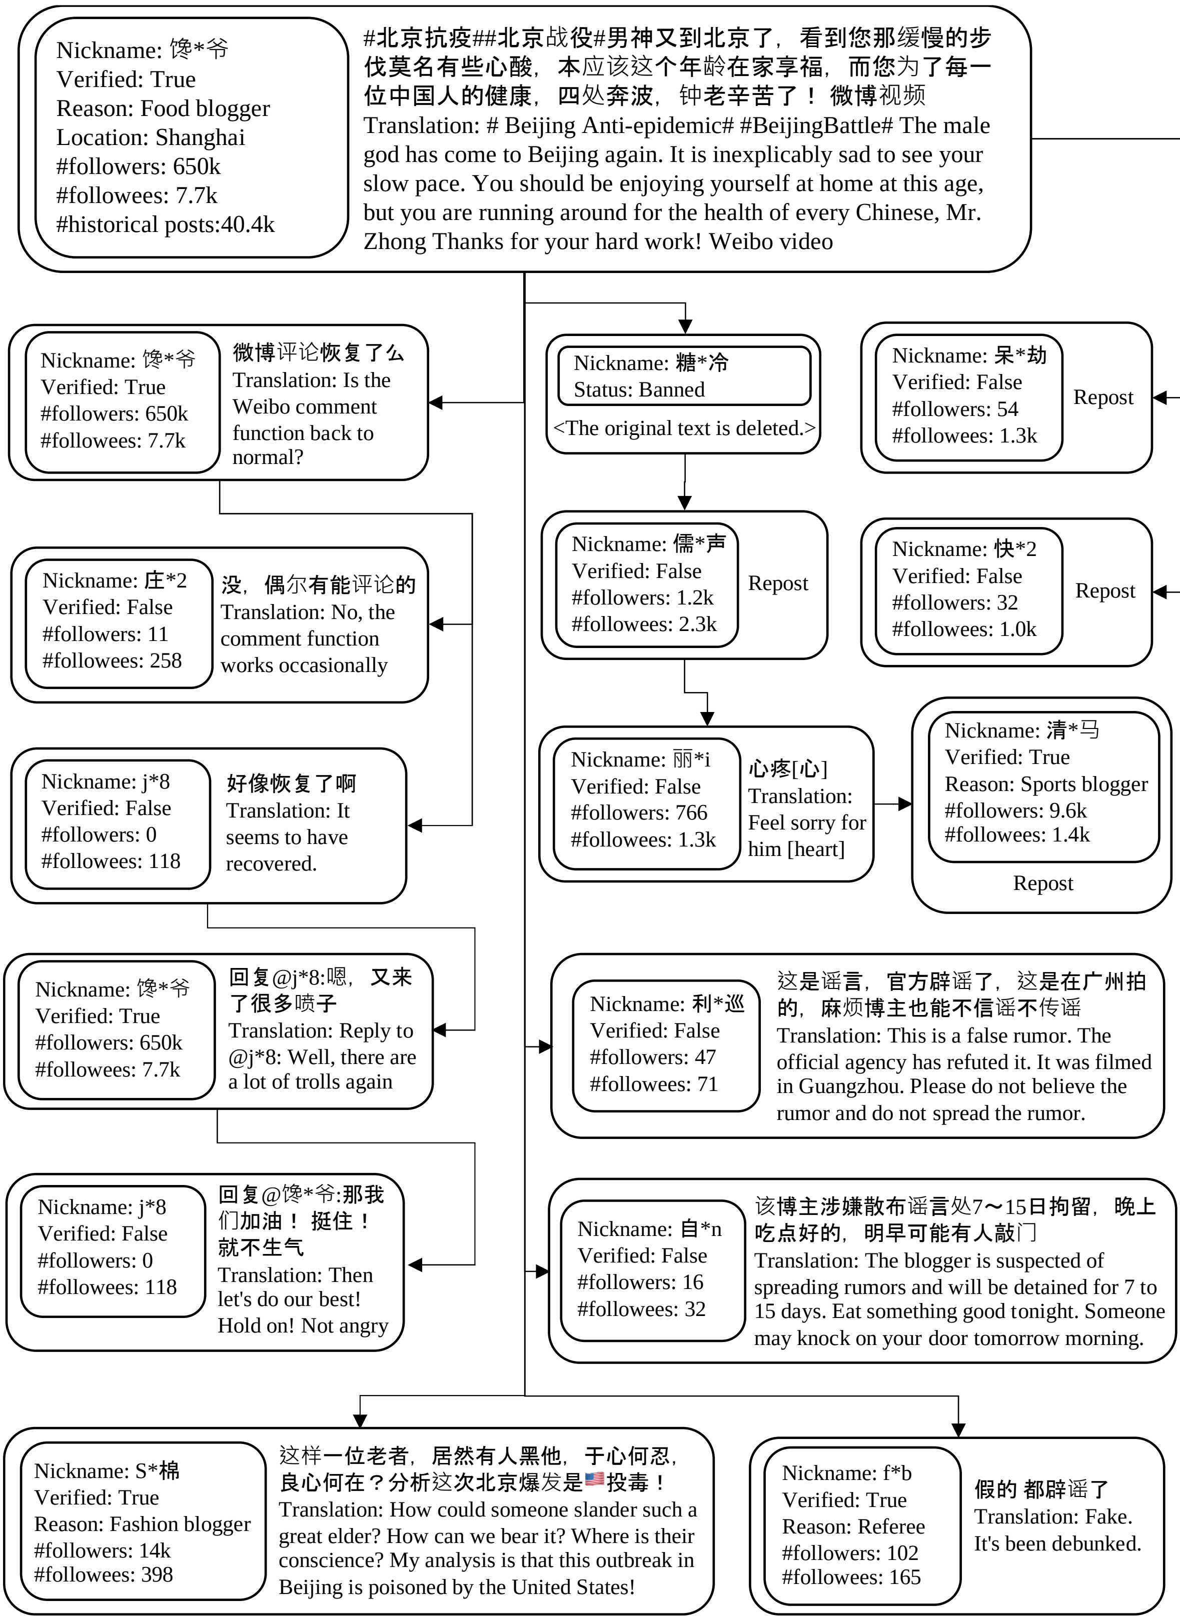

Figure 3.1. An example of cascade structure

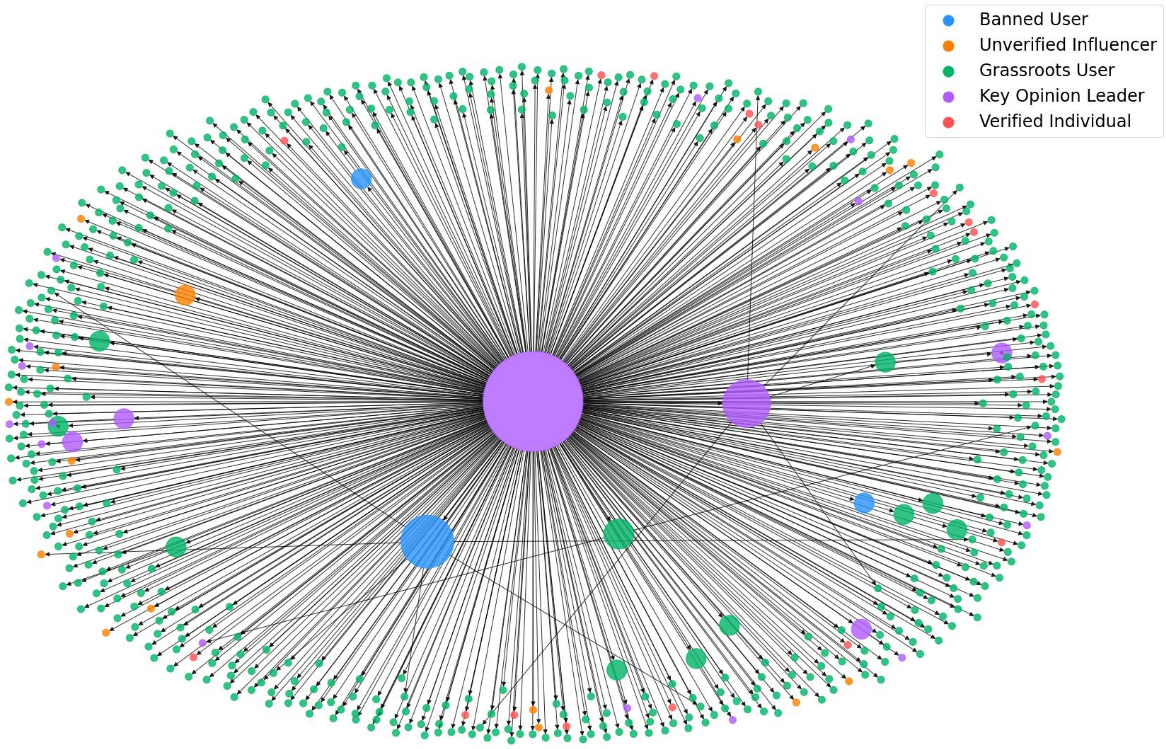

Figure 3.2. A complete cascade of misinformation diffusion

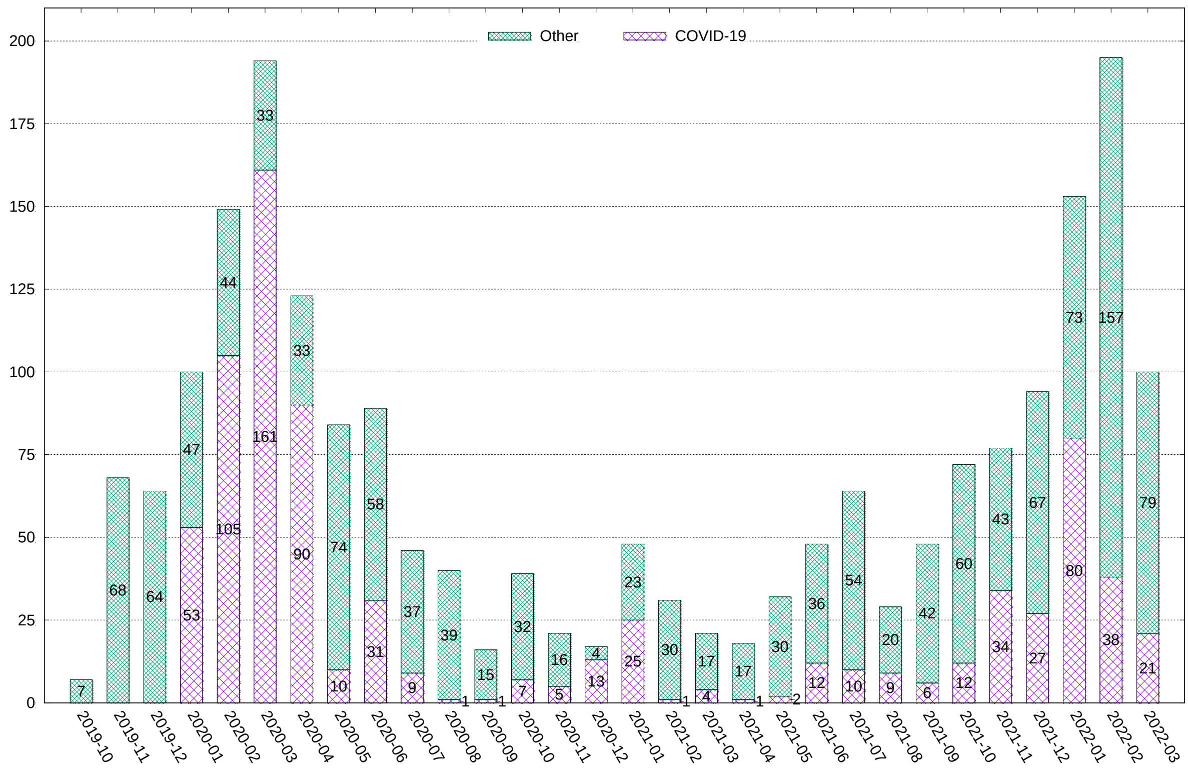

Figure 3.3. Misinformation over time

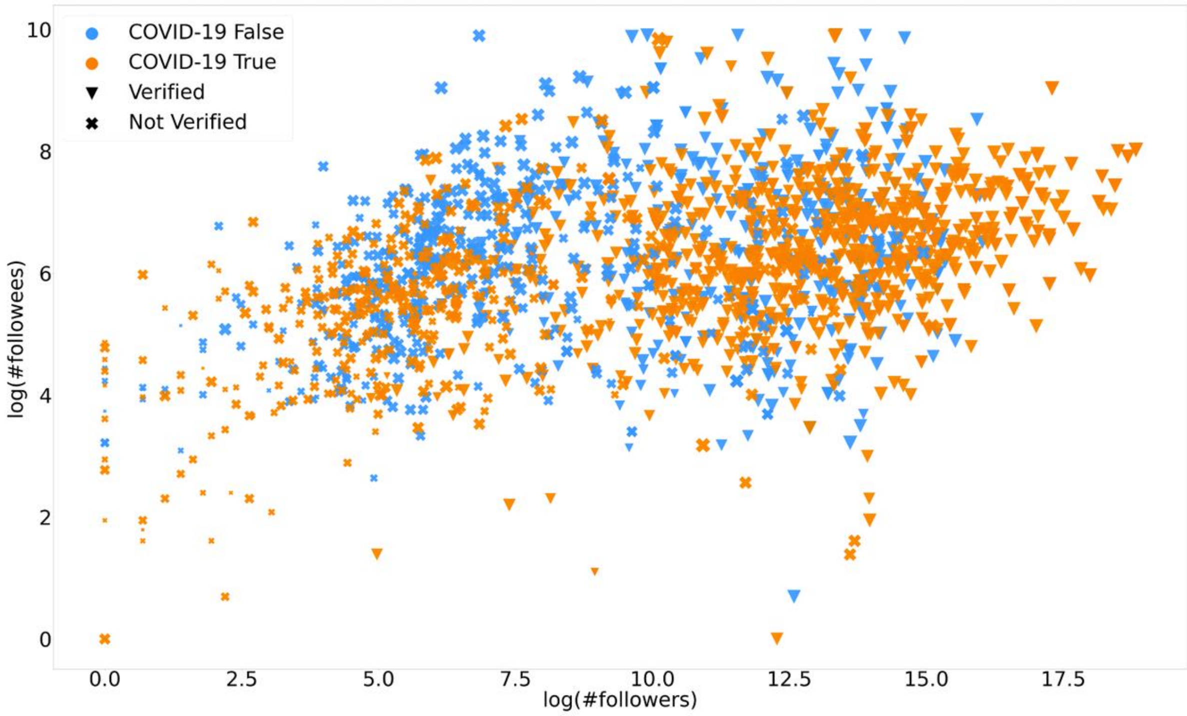

Figure 3.4(a). User profile information of source publishers where areas of points represent the number of historical posts

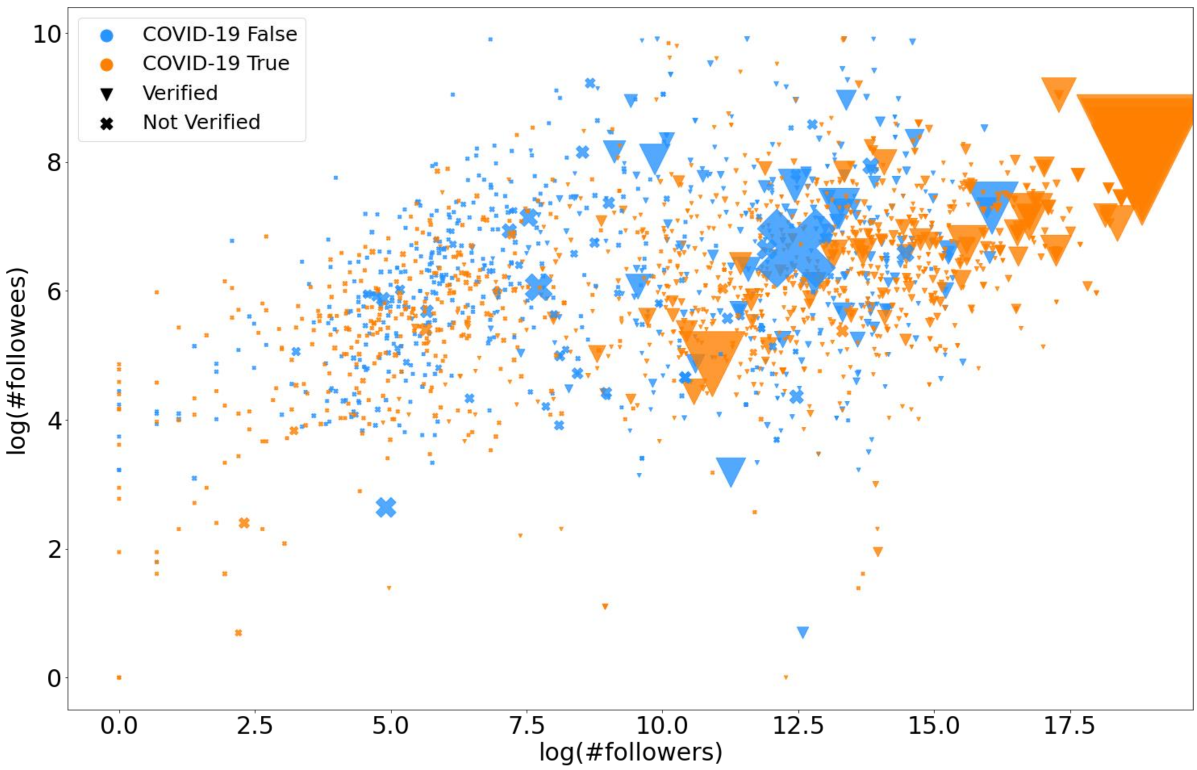

Figure 3.4(b). User profile information of source publishers where areas of points represent cascade size

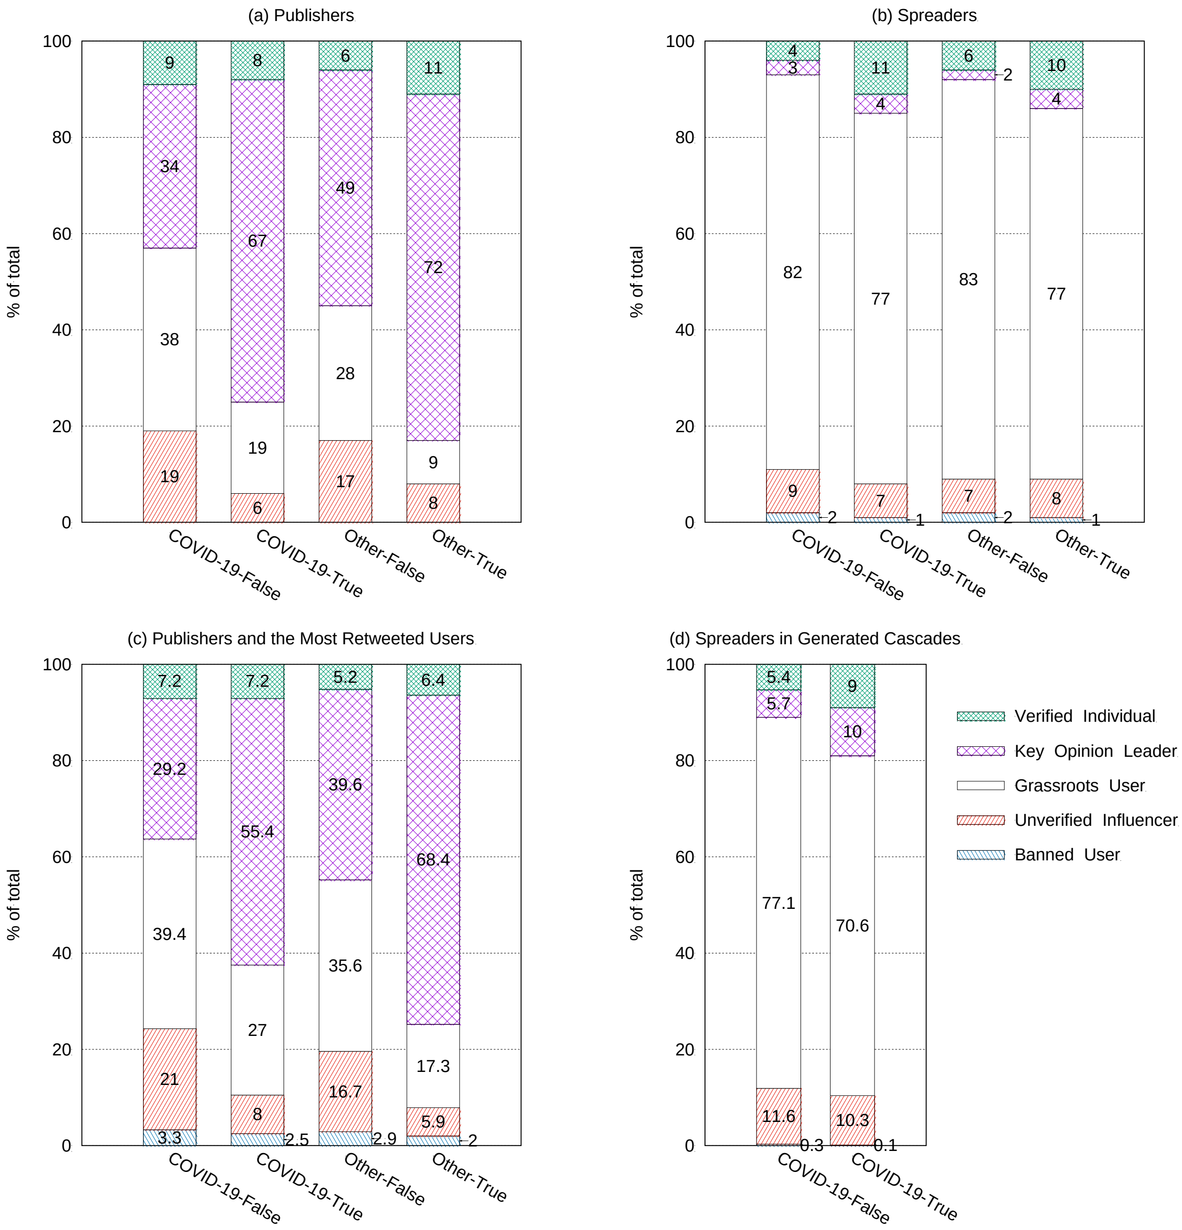

Figure 3.5. User type distributions under different conditions