Hindawi
Security and Communication Networks
Volume 2021, Article ID 6656470, 12 pages
https://doi.org/10.1155/2021/6656470

*Research Article*

# Zombie Follower Recognition Based on Industrial Chain Feature Analysis

**Juan Tang**,[1] **Hualu Xu**,[2] **Pengsen Cheng**,[2] **Jiayong Liu**,[2] **Cheng Huang**,[2] **and Xun Tang**[2]

[1]*College of Electronics and Information Engineering, Sichuan University, Chengdu 610065, China*
[2]*School of Cyber Science and Engineering, Sichuan University, Chengdu 610065, China*

Correspondence should be addressed to Jiayong Liu; ljy@scu.edu.cn

Received 20 November 2020; Revised 3 January 2021; Accepted 25 January 2021; Published 8 February 2021

Academic Editor: Zhe-Li Liu

Zombie followers, a type of bot, are longstanding entities in Sina Weibo. Although the features and detection of zombie followers have been extensively studied, zombie followers are continuously increasing in social networks and gradually developing into a large-scale industry. In this study, we analyze the features of eight groups of zombie followers from different companies. The findings indicate that although zombie followers controlled by different companies vary greatly, some industries may be controlled by the same organization. Based on the feature analysis, we use multiple machine learning methods to detect zombie followers, and the results show that zombie follower groups with short registration time are more easily detected. The detection accuracy of zombie followers that have been cultivated for a long duration is low. Moreover, the richer the feature sets, the higher the recall, precision, and $F_1$ of their detection results will be. Under a given rich feature set, the accuracy of the combined-group detection is not as high as that of the single-group detection. The random forest achieves the highest accuracy in both single- and combined-group detections, yielding 99.14% accuracy in the latter case.

## 1. Introduction

Sina Weibo is an online social network service, such as Twitter and Facebook, and has nearly 516 million monthly active users till December 2019 according to the fourth quarter financial report of Sina Weibo [1]. Similar to the case of other social media services, many misbehaving accounts [2] exist in Sina Weibo, such as bots [3–5], trolls [6–8], sockpuppets [9, 10], and compromised accounts [11, 12]. The ultimate aim of such accounts that participate in social networks is to cause disruption of the normal order.

Zombie followers [13–15], a type of bot [3, 4, 16], are longstanding entities in Sina Weibo. They are often used to spread malicious information, manipulate public opinion, steal personal information, and so on [4, 17–19]. They not only undermine users' social credibility but also adversely affect users' network security and social environment [15, 20, 21]. Researchers in the past have often focused on the detection of zombie followers [15, 22]. They analyzed the feature differences between zombie followers and normal users, such as text features [23, 24], behavior features [25, 26], or network structure features [27–29], and then combined machine learning methods for zombie follower detection [30, 31].

Although the features and detection of zombie followers have been extensively studied, zombie followers are continuously increasing in online social networks and gradually developing into an industry [26, 32]. We observed that zombie followers on Sina Weibo are gradually moving toward this trend and forming a large-scale ecosystem, wherein the user can get many zombie followers for a small cost. Previous studies have not analyzed the characteristics of different zombie follower industries. Questions such as will there be differences in the characteristics of zombie followers from different sources and in the detection results if the same detection method is used for zombie followers with different characteristics have not been explored. Therefore, the study of zombie followers' ecosystem and industry features will help us better understand and automatically detect them.

*Present work*. In this paper, we focus on the feature analysis and detection of eight zombie follower groups. Herein, zombie followers [15] are defined as malicious users that are manipulated and maintained by programs. They imitate human behaviors and influence normal social behaviors on social networks. The zombie follower industry [33] is defined as a new black market formed by merchants engaged in the production and sale of zombie followers. To analyze the characteristics of the zombie follower industry, we investigate various organizations engaged in the trading of zombie followers on the Internet. Here, an organization or company that provides sales of zombie followers is termed a zombie follower company (hereafter, ZF company).

We collected eight zombie follower groups (each group having more than 5,000 accounts) from different ZF companies. Based on the collected data, our paper provides the following three main contributions:

(1) We analyzed the basic features and content features of zombie follower groups and found that zombie follower companies always mass produce zombie followers. Due to varying registration time and service scope, those zombie followers usually have different features.

(2) We study the interactive relationship between the eight zombie follower groups. The findings indicate that some of the merchants selling zombie followers are actually operated by the same organization.

(3) Finally, based on the study of the aforementioned features of the zombie follower industry, we use machine learning methods to detect the zombie followers in a single group and in combination. In the single-group detection, zombie followers with short registration time are more easily detected. The detection accuracy of zombie followers that have been cultivated for a long time declines. Moreover, the richness of the feature set plays an important role in the detection. The richer the feature sets, the higher the recall, precision, and $F_1$ values will be. Although the accuracy of the combined-group detection is not as high as that of the single-group detection, the random forest is the highest in both detections, with 99.14% accuracy in the combined-group detection.

## 2. Data Specification

This section details the source of our dataset as well as the settings of honeypot accounts and the data crawling process.

*2.1. Data Source.* On Weibo, the number of followers of users often depends on the users' influence. Driven by the benefits of pan-entertainment and commercialization, users' demand for zombie followers has grown, leading to a large-scale purchase of zombie followers in the market. Based on the investigation of the various advertisements on the Internet to sell zombie follower services, the following main sales channels can be found:

Weibo profile: some zombie followers leave sales advertisements with contact information on normal users' microblogs, while others mark such information on their avatars and spread it by following normal users.

Taobao shop: Taobao, a popular C2C platform in China, has taken all efforts to stop illegal sales, but a search using specific keywords can still lead one to the sales of zombie follower services. The stores offer various packages and for each package, the basic information, quantity, and price of zombie followers are explained on the product details page. Buyers can place orders directly according to the instructions.

Search engine: when searching for keywords related to the sales of zombie follower services on major search engines, a series of related websites will appear, which contain information such as the categories and number of zombie followers. Buyers can purchase directly from the website or contact the customer service staff according to the information provided on the website.

*2.2. Honeypot Account Description.* Honeypot is a common means to collect zombie followers [32, 34]. On Facebook, MySpace, and Twitter, it is often used to detect spammers [35–37]. Aiming at studying the current ecosystem of zombie followers, we created eight honeypot accounts on Weibo, corresponding to the eight ZF company groups.

All honeypots remain in the initial state and empty (i.e., no basic information or microblog is present). We collected more than 5,000 zombie followers each from eight different companies and injected them into the corresponding honeypot account. Table 1 lists honeypot account details and the sources of zombie followers. All zombie followers were collected at the same time. Overall, we collected a total of 43,352 zombie followers in eight groups.

*2.3. Data Collection.* What are the characteristics of the zombie follower industry? What are the characteristic differences between zombie followers and normal users? How to detect zombie followers? To answer these questions, in this study, we mainly collected two datasets: (1) the zombie followers' data and related data collected through the honeypot account and (2) the normal users' data and related data collected through the Python crawler. All the above data were exclusively open data obtained through Sina API. In addition, we encrypted the data to ensure data security. The collection of the datasets and the detailed analysis of the basic characteristics are described below.

We used the Python crawler to monitor the corresponding honeypot account and detect the injected zombie followers. After all zombie followers were collected, we performed a second crawl on the collected data. The public information, followers (the latest 1000), and microblogs (the latest 100) of each zombie follower were crawled, and the results are presented in Table 2. Meanwhile, we collected the data, including public information and 3,394,129 microblogs

Table 1: Summary statistics of honeypot accounts.

| User ID | Campaign name | Provider/source | #Zombie followers |
|---|---|---|---|
| 01_72****37 | A000 Douyin and Weibo flagship shop | Weibo profile | 5069 |
| 02_72****04 | Sihui Network | Taobao shop | 5270 |
| 03_72****91 | Xingchen Network Technology | Search engine | 6172 |
| 04_72****62 | Aijia Network | Search engine | 5258 |
| 05_72****21 | Self-help business platforms for Douyin, Kuaishou, and Weibo | Search engine | 6313 |
| 06_72****43 | A Weibo-Douyin-WeChat platform | Weibo profile | 5091 |
| 07_72****47 | Niuweifen marketing | Taobao shop | 5073 |
| 08_72****28 | Yunyidingdian platform | Weibo profile | 5106 |

Table 2: Summary statistics of the eight honeypot accounts.

| User ID | #Followers | #Followers obtained | #F_followers[1] | #Microblogs |
|---|---|---|---|---|
| 01_72****37 | 5,069 | 5,063 | 661 | 20,528 |
| 02_72****04 | 5,270 | 943 | 21 | 3 |
| 03_72****91 | 6,172 | 6,145 | 227,101 | 285,756 |
| 04_72****62 | 5,258 | 5,257 | 79 | — |
| 05_72****21 | 6,313 | 5,000 | 80 | — |
| 06_72****43 | 5,091 | 5,022 | 1,081,774 | 642,679 |
| 07_72****47 | 5,073 | 5,072 | 2,550 | 0 |
| 08_72****28 | 5,106 | 5,105 | 1,186 | 18,584 |

[1] F_followers are the followers of the zombie followers in the eight zombie follower groups.

(the latest 100) of 45,559 normal users as a comparison dataset.

Each ZF company assures us that their products are authentic and reliable. Their zombie followers have avatars, personal information, and irregular updates of microblogs. More importantly, they cannot be blocked. However, the zombie followers of some ZF companies were blocked by Sina Weibo within a short time. For example, Groups 02, 04, and 05 were blocked shortly after the infusion was completed. Among them, 943 zombie followers belonging to Group 02 were unblocked after some time, but they were soon blocked again. Groups 04 and 05 were banned in the early period, so some of their data were missing (only Weibo ID, user name, and the number of followers and friends could be collected). In addition, we found that a small number of zombie followers were blocked in other groups. Finally, we obtained a total of 37,607 zombie followers, having 1,313,452 followers and 967,550 microblogs.

## 3. Analysis of Characteristics

In previous studies, zombie followers and normal users were usually distinguished from various perspectives [15], such as users' personal information [13, 14], relationship features [28], behavioral features [26], and emotional features. In this section, we attempt to answer the following two questions on the basis of basic features and content features: what is the difference between zombie follower groups and what is the difference between zombie follower groups and the normal user group?

*3.1. Basic Characteristics.* We randomly selected 5000 users from the normal user group as Group 09. For comparison, we also calculated the average value of Groups 01–08 and considered it as Group 00. In this section, we combine existing fields in the dataset and compare the groups in terms of five aspects: registration time, the number of followers and friends of users, username complexity, and user hierarchy.

*3.1.1. More Centralized Registration Time.* Figure 1 shows the cumulative distribution function (CDF) graphs of registration time of users in all the groups. It indicates that the registration time of the normal user group (Group 09) is evenly distributed. However, the CDF graphs of the zombie follower groups (data in Groups 02, 04, and 05 were missing) are significantly different from that of Group 09. The distribution graphs of Groups 01 and 08 are similar, with a stepped increase and high consistency in value. Groups 03 and 06 are also similar, and their graphs are closer to that of Group 09. However, compared with the graph of Group 09, graphs of Groups 03 and 06 are not smooth and show a stepped increase. Their distribution is similar to that of Groups 01 and 08. All the zombie followers in Group 07 were registered three days before purchase, so its CDF is more concentrated. As Figure 1 shows, most of the zombie followers were produced recently. Therefore, we can infer that most ZF companies continue to mass produce zombie followers.

In summary, zombie follower groups manipulated by different ZF companies have significant differences in registration time. Some ZF companies hold and mass produce zombie followers close to the purchase time, while others mass produce them in advance. As a comprehensive CDF graph of zombie followers, the curve of Group 00 indicates that the zombie follower industry is developing on a large scale.

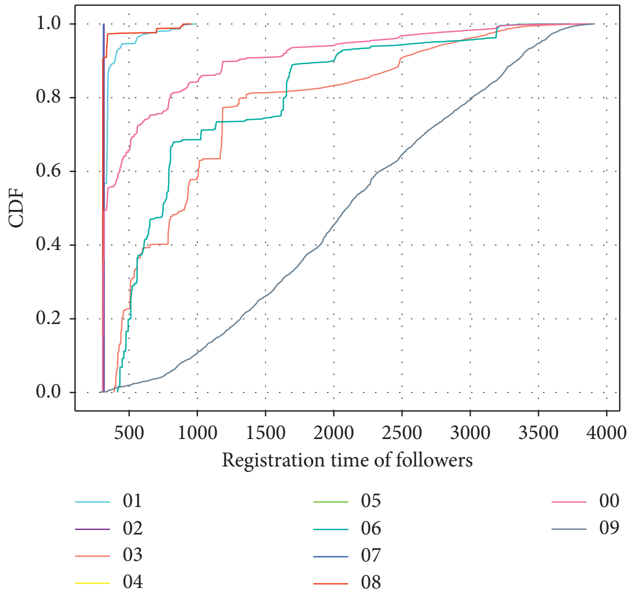

Figure 1: Cumulative distribution of registration time of users in different groups (00 is the average registration time of zombie follower groups; 09 is the registration time of the normal user group).

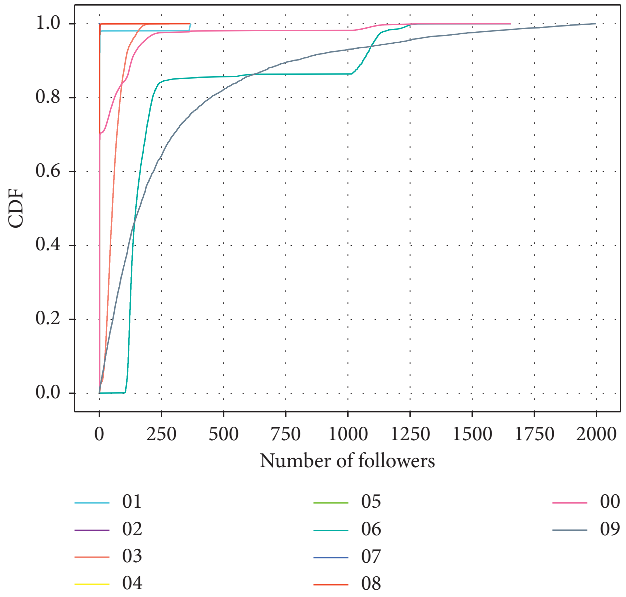

Figure 2: Cumulative distribution of the number of user followers in different groups.

### 3.1.2. *Fewer Followers and Mutually Following.*

In Figure 2, the median number of user followers in Group 09 is 186, whereas that of Groups 01–08 is 29, indicating a large difference between them. Most zombie followers of Groups 04, 05, and 07 have only one to three followers. The distribution of Groups 01 and 08 is similar, and most of their users have no followers. The followers in Group 03 are relatively dispersed, evenly distributed between 0 and 250, while those in Group 06 are almost all over 100 (only two users have less than 100).

Based on the above results, compared with Group 09, the number of user followers in Groups 01–08 is more consistent. A cursory investigation reveals that to make zombie followers resemble normal users, zombie followers in Groups 01–08 generally follow each other, thus forming a network of zombie followers.

### 3.1.3. *Prefer to Be a Follower Based on the Service.*

As shown in Figure 3, the distribution of the number of friends of users in Group 09 is even, with about 80% being less than 500. By contrast, the distribution of Groups 01–08 is irregular and most of the zombie followers have more friends than normal users have, such as Groups 03 and 06. The registration time distribution for Groups 03 and 06 suggests that they have been engaged in selling follower services for a long duration, so their users have more friends. Furthermore, most users in Groups 01 and 08 have less than 200 friends. In Groups 04, 05, and 07, all have fewer than 200 friends, and their CDF graphs show an irregular stepped increase (Figure 4). Based on the registration time distribution, we believe that most zombie followers in the five groups have not been engaged in the business for a long duration, so they have not accumulated a large number of friends.

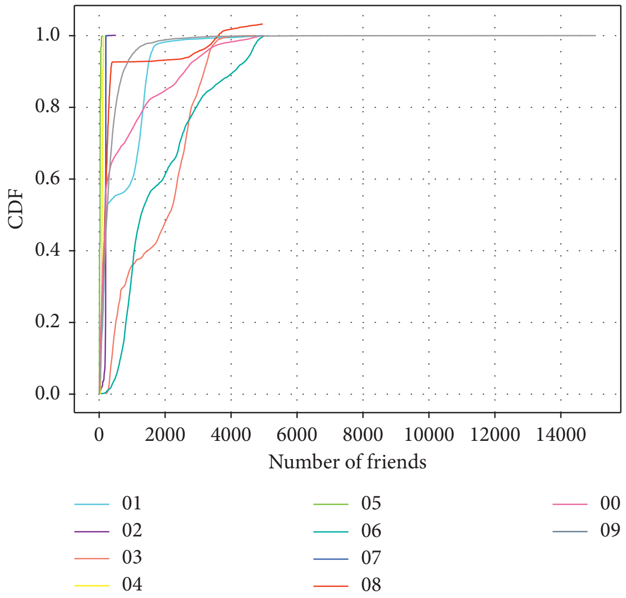

Figure 3: Cumulative distribution of the number of friends in different groups.

To better reflect the composition of users' social relations, we use the interaction index function [38], defined as

$$\text{interaction index} = \frac{\text{followers count}}{\text{friends count}}. \tag{1}$$

Figure 5 shows the CDF graphs of the interaction index for all groups. In Group 09, the interaction index of 96.76% of the users is less than 10, and the maximum index is 81. In Groups 01–08, the interaction index of only 16.09% is less than 10, while that of 26.12% is greater than 81.

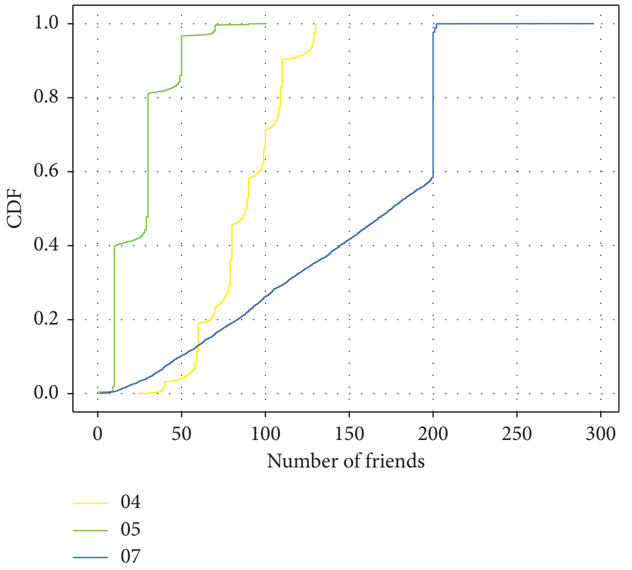

Figure 4: Cumulative distribution of the number of friends of users in Groups 04, 05, and 07; an irregular ladder increase in the number is observed.

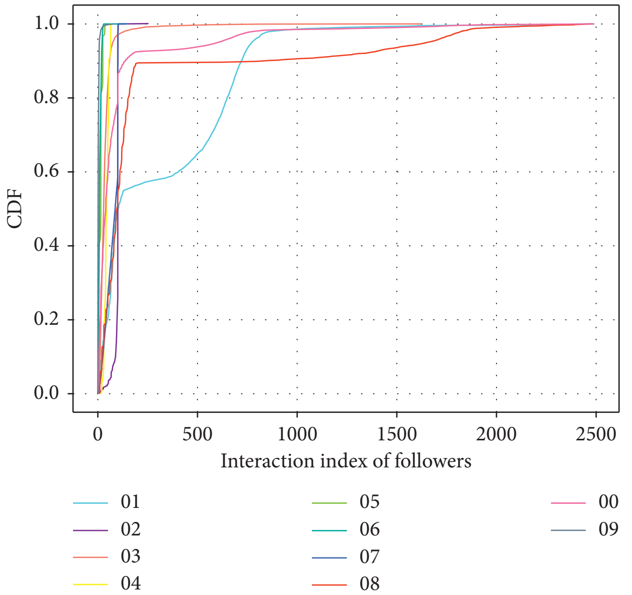

Figure 5: Cumulative distribution of the interaction index for different groups.

*3.1.4. Simpler and Meaningless Usernames.* We next analyze the username complexity of Groups 01–09 by using the Jieba [39] algorithm to segment the usernames. Accordingly, if $n$ is the number of words in the username, $K$ is the number of numerals, and $\mathrm{len}_i$ is the length of the $i$-th word, then the complexity of the username [40] is given as

$$\text{complex} = n + \sum_{i=1}^{k} \frac{\mathrm{len}_i}{3}. \quad (2)$$

The username complexity of Groups 01–09 is shown in Figure 6. The figure indicates that the username complexity of Group 09 is greater than that of Groups 01–08. The analysis of the composition of usernames in Groups 01–08 reveals that the usernames of Groups 04 and 05 are automatically generated by the system, with their structure being "user" + random number. Compared with other groups, Groups 01 and 08 usernames are more readable and have specific rules for their generation. Most usernames in Groups 02, 03, 06, and 07 are random combinations of Chinese characters and letters, bearing no specific meaning.

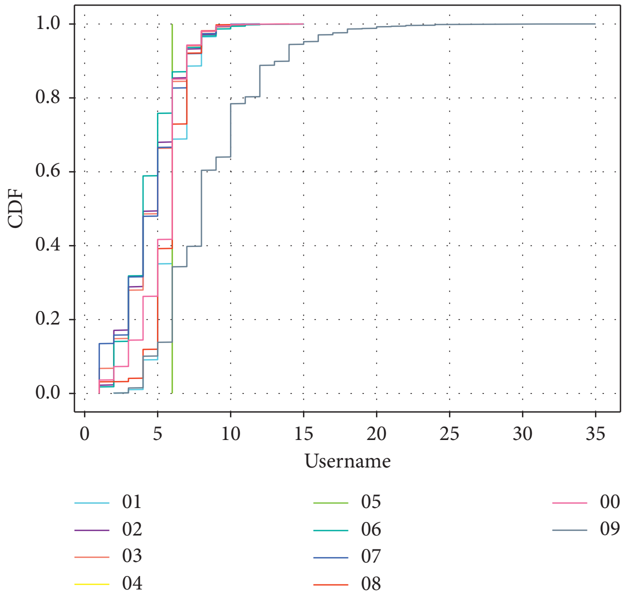

Figure 6: Cumulative distribution of the username complexity in different groups.

*3.1.5. Lower Level in Weibo's Hierarchy.* Our calculation of the hierarchies of users in Groups 01–09 shows that out of the 43,352 zombie followers, only 1,278 have a user hierarchy greater than 0, accounting for 4.84%. However, in Group 09, 2945 users are greater than 0, accounting for 58.9%. Due to the mass production of zombie followers, ZF companies cannot improve the hierarchies for most zombie followers. However, it is worth noting that ZF companies not only sell zombie follower services but also control mass social robots with advanced authentication and higher hierarchies than normal users. In the future, we will conduct research considering this aspect.

*3.2. Content Characteristics.* This section presents a comparison of the relevant features of users' microblogs in Groups 01–09. We could not obtain the microblog data of Groups 04 and 05 because the zombie followers in these groups were blocked. Moreover, 5072 zombie followers in Group 07 did not post any microblog as their registration time was shorter, and 943 zombie followers in Group 02 only posted three microblogs. Therefore, we were focusing on the content of only Groups 01, 03, 06, 08, and 09.

*3.2.1. Replication of Original Content.* To determine what content zombie followers often post, we analyzed the microblog content of Groups 01–09. As Table 3 shows, every microblog in Groups 01 and 08 is a repost, and mostly the same posts. Conversely, Groups 03 and 06 are more balanced, including original microblogs and reposts. Among them, most of the original microblogs repeat celebrity quotes or common senses. The repetition rate of reposts of zombie followers is generally higher than that of normal users. Moreover, in Groups 03 and 06, the rate is 50% or more. Thus, we can infer that the zombie followers manipulated by ZF companies hardly post original microblogs and their reposts are related to their business. Therefore, ZF companies are suspected of manipulating public opinion.

*3.2.2. Low Interaction of Microblogs.* By analyzing the content of microblogs of Groups 01–09, we investigated if zombie followers write differently from normal users. As Table 4 shows, compared to Group 09, zombie follower groups contain fewer URLs, mentions (@), and hashtags, and they are less interactive with other users. In addition, although the length of microblogs of different groups is different, groups with similar basic characteristics have similar length of microblogs.

*3.2.3. Poor Microblog Sources.* Regarding the microblog sources, Table 5 shows that Group 09 has 3,650 sources, accounting for 94.32% of the total (3,870), whereas the sources of zombie follower groups are considerably less. Among them, Groups 01 and 08 have only three consistent sources, whereas Group 06 has the most abundant sources (only 429).

*3.2.4. Poor Spreading of Microblogs.* The communication features [26] of the microblogs of Groups 01–09 are analyzed in Table 6. As the table indicates, the zombie follower groups are significantly different from Group 09 as more than 98% of the zombie followers have zero reposts, attitudes, and comments, and almost no group has a count above 10. We conclude that although zombie followers do post microblogs, they usually get little attention from other users; thus, the posts have poor ability to spread.

*3.3. Discussion.* Zombie follower groups have different features because of varying registration time. Due to longer survival time, Groups 03 and 06 have greater similarity to normal users. However, the registration time of Group 07 is only three days, so all its features are notably different. It can be expected that some ZF companies have been engaging in cultivating zombie followers for a long duration, and the longer they manipulate, the more similar the zombie followers will be to normal users.

Comparison of the features of zombie follower groups reveals an interesting phenomenon. The features of Groups 04 and 05, Groups 01 and 08, and Groups 03 and 06 are very similar. In the next section, we will analyze whether these zombie follower groups are correlated.

## 4. Ecosystem Characteristics

We focus here on the following three questions: why is it difficult for normal users to identify zombie followers? Are different ZF companies correlated? Are there social relations among zombie follower groups?

*4.1. Why Is It Difficult for Normal Users to Identify Zombie Followers?.* From the features of zombie follower groups described in Section 3, we can conclude that zombie followers are considerably different from normal users. However, it is often difficult for normal users to judge whether a user is a zombie follower. Note that when normal users visit others' profiles, they usually judge the profile authenticity by observing its basic information and microblogs.

We analyzed the differences in the basic information between Groups 01–08 and Group 09 (Table 7). Compared with normal users, zombie followers (95.83%) have more complete basic information. To avoid blocking of the zombie followers, ZF companies make them behave more like humans. For example, Groups 03 and 06 not only have extremely complete basic information (99.86%) but also have rich original microblogs (Table 4). Clearly, zombie followers also have real avatars, complete basic information, and simulated original microblogs and reposts. Therefore, the boundary between the zombie followers and normal users becomes increasingly blurred, making distinction of zombie followers difficult.

*4.2. Are Different ZF Companies Correlated?.* We analyzed the relationship among zombie follower groups to determine if ZF companies are correlated, if these companies manipulate different zombie follower groups with different sales methods and if these companies belong to the same organization.

By observing the interaction of each user from Groups 01–08, we determine if any zombie follower is present in two or more groups simultaneously. After matching, we found 8 identical zombie followers in Groups 01 and 08 and 104 in Groups 04 and 05 (Figure 7). Considering the characteristics of these groups, it is reasonable to conclude that they are controlled by the same organization.

*4.3. Are There Social Relations among Zombie Follower Groups?* After determining that there may be some correlation among zombie follower groups, we attempt to determine whether there also have some correlating social relations.

The analysis of the followers of zombie followers indicates that most of them are zombie followers. Therefore, we compared the followers of the zombie followers in Groups 01–08. The results demonstrate a small amount of overlap between some groups (Table 8). However, the number of overlaps in Groups 03 and 06 is as high as 5,696. After removing the repeated followers, we obtained 218 followers of the zombie followers in Groups 03 and 06 multiple times.

Table 3: Summary statistics of microblog content.

| Collection | Original count | Repost count | Repetition | Repetition rate (%) |
|---|---|---|---|---|
| 01 | 0 | 20,528 (100%)[1] | 2,552 | 31.41 |
| 03 | 252,690 (88.43%) | 33,066 (11.57%) | 5,493 | 65.41 |
| 06 | 312,630 (48.64%) | 330,049 (51.36%) | 44,652 | 55.39 |
| 08 | 0 | 18,584 (100%) | 2,400 | 31.46 |
| 09 | 145,768 (37.05%) | 247,720 (62.95%) | 23,506 | 23.73 |

[1] Parenthetical information is the percentage of the total number of microblogs in each group.

Table 4: Summary statistics of the microblog content features.

| Collection | Status count | URLs | Mentions | Hashtags | Ave-length |
|---|---|---|---|---|---|
| 01 | 20,528 | 76 (0.37%)[1] | 996 (4.85%) | 0 | 12.42 |
| 03 | 285,756 | 10,620 (3.72%) | 6,215 (2.17%) | 19,225 (6.73%) | 141.99 |
| 06 | 642,679 | 9,821 (1.53%) | 66,379 (10.33%) | 66,534 (10.35%) | 65.94 |
| 08 | 18,584 | 60 (0.32%) | 751 (4.04%) | 0 | 11.70 |
| 09 | 393,488 | 85949 (21.84%) | 76,682 (19.49%) | 202,010 (51.34%) | 55.56 |

[1] Parenthetical information is the percentage of the total number of microblogs in each group.

Table 5: Summary statistics of the microblog sources.

| Collection | 01 | 03 | 06 | 08 | 09 |
|---|---|---|---|---|---|
| Count | 3 (0.08%) | 429 (11.09%) | 217 (5.61%) | 3 (0.08%) | 3,650 (94.32%) |

Table 6: Summary statistics of the microblog communication features.

| Collection | Range | Reposts (%) | Attitudes (%) | Comments |
|---|---|---|---|---|
| 01 | 0 | 99.58 | 98.82 | 99.92% |
| | 0+ | 0.42 | 0.18 | 0.08% |
| 03 | 0 | 99.79 | 97.07 | 99.47% |
| | 1–10 | 0.19 | 2.81 | 0.51% |
| | 10+ | 0.02 | 0.12 | 0.02% |
| 06 | 0 | 99.61 | 98.12 | 99.61% |
| | 1–10 | 0.38 | 1.87 | 0.39% |
| | 10+ | 0.01 | 0.01 | 0 |
| 08 | 0 | 99.73 | 98.99 | 99.88% |
| | 1–10 | 0.27 | 1.01 | 0.12% |
| 09 | 0 | 92.27 | 71.72 | 82.11 |
| | 1–10 | 5.60 | 23.31 | 14.03% |
| | 10+ | 2.13 | 4.97 | 3.86% |

Although there are no identical zombie followers in Groups 03 and 06, their potential social relations suggest that they may belong to the same organization.

## 5. Detection Model

*5.1. Experimental Features.* As described in the above sections, we studied the behavior of zombie followers from the following three aspects and list the relevant features in Table 9: (1) basic features, which include the complexity of the user name, the number of followers and friends of users, the interaction index, hierarchy, and registration time of users. (2) Content features, including the rate of original microblogs and reposts, the total number of microblogs of users, URLs, hashtags, and mentions, and the average length of all microblogs. (3) Ecosystem features, it refers to the user's gender and age and whether the basic information of the user is provided.

*5.2. Experimental Design.* In this study, Python was used to realize the whole process of feature extraction and model construction, as shown in Figure 8. This model mainly consists of two parts: the data feature analysis module and zombie follower detection module. In the data feature analysis module, we obtain the original data set through the crawler, remove the invalid data after preprocessing, and format them. We then analyze the data features and transform them into the corresponding feature set. In the zombie follower detection module, we use five-fold cross-validation. The detailed data distribution is shown in Table 10. Finally, three machine learning methods (KNN [41], SVM [42], and random forest [43]) are used to detect the

Table 7: Summary statistics of the basic info of zombie followers.

| Collection | Avatar | Location (%) | Gender (%) | Age (%) | Simple info (%) |
|---|---|---|---|---|---|
| 01 | No | 77.07 | 77.64 | 40.61 | 77.68 |
| 02 | Yes[1] | 99.26 | 100 | 93.96 | 100 |
| 03 | Yes | 80.47 | 98.34 | 51.91 | 98.34 |
| 06 | No | 84.73 | 99.88 | 9.66 | 99.88 |
| 07 | Yes | 98.68 | 99.82 | 93.24 | 99.82 |
| 08 | No | 88.18 | 99.24 | 51.23 | 99.24 |
| 09 | Yes | 54.36 | 77.56 | 38.72 | 77.56 |

[1] "Yes" indicates that most of the zombie followers have avatars.

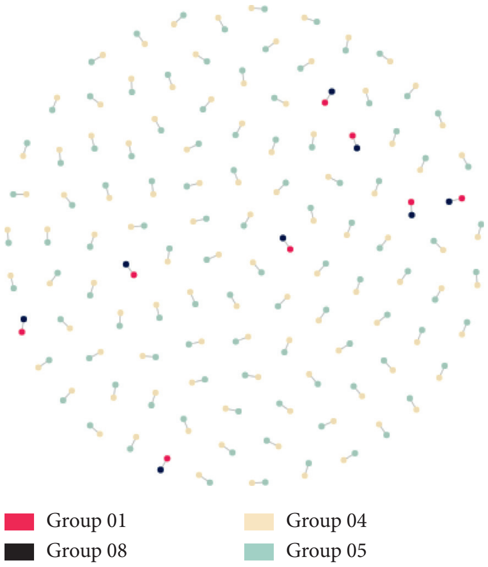

Figure 7: Matching of identical zombie followers in Groups 01–08.

Table 8: Summary statistics of the overlap of the followers of zombie followers.

| Group | Num |
|---|---|
| 01 and 03 | 10 |
| 01 and 06 | 20 |
| 01 and 08 | 18 |
| 03 and 04 | 2 |
| 03 and 05 | 2 |
| 03 and 06 | 5696 |
| 03 and 07 | 6 |
| 03 and 08 | 7 |
| 04 and 05 | 8 |
| 06 and 07 | 90 |
| 06 and 08 | 6 |
| 01, 03, and 06[1] | 4 |
| 01, 03, and 08 | 2 |
| 03, 04, and 05 | 2 |
| 01, 03, 06, and 08 | 1 |

[1] Only the overlapped groups were output.

experimental data. Each of the three algorithms uses default parameters, and the calculation is executed on a computer with Intel(R) Xeon(R) CPU and 8 GB memory.

We evaluated the performance of the proposed model by using the indicators accuracy, recall, precision, and $F_1$. They are defined as follows:

$$\begin{aligned}\text{Accuracy} &= \frac{\text{TP} + \text{TN}}{\text{TP} + \text{TN} + \text{FT} + \text{FN}},\\ \text{Recall} &= \frac{\text{TP}}{\text{TP} + \text{FN}},\\ \text{Precision} &= \frac{\text{TP}}{\text{TP} + \text{FP}},\\ F_1 &= \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}},\end{aligned} \tag{3}$$

where TP (true positive) is the number of normal users predicted as normal users; FP (false positive) is the number of zombie followers predicted as normal users; TN (true negative) is the number of zombie followers predicted as zombie followers; and FN (false negative) is the number of normal users predicted as zombie followers.

*5.3. Experimental Results.* In our experiments, we attempted to maximize the detection of zombie follower groups. We extracted the features of each user described in the previous sections and used three types of classifiers (i.e., KNN, SVM, and random forest) to detect eight zombie follower groups. Here, when the detection method uses any set of data from 01 to 08 zombie follower data and normal user data, it is called the single-group detection, and when it uses eight sets of zombie user data and normal user data, it is called the combined-group detection. Meanwhile, we also detected basic features, content features, and ecosystem features in single-group and combined-group detections, and the results are shown in Tables 11–13.

In the single-group detection, when part of data is missing (Groups 04 and 05 lack in content features and ecosystem features; Groups 02 and 07 lack in content

Table 9: Description of feature classifications.

| Feature set | Features |
|---|---|
| Basic features | Username complexity, number of followers and friends, interaction index, user hierarchy, user registration time |
| Content features | Number of original post, repost, and status; number of URLs, hashtags, mentions, average length of microblogs (Avg-length) |
| Ecosystem features | Gender, age, simple info |

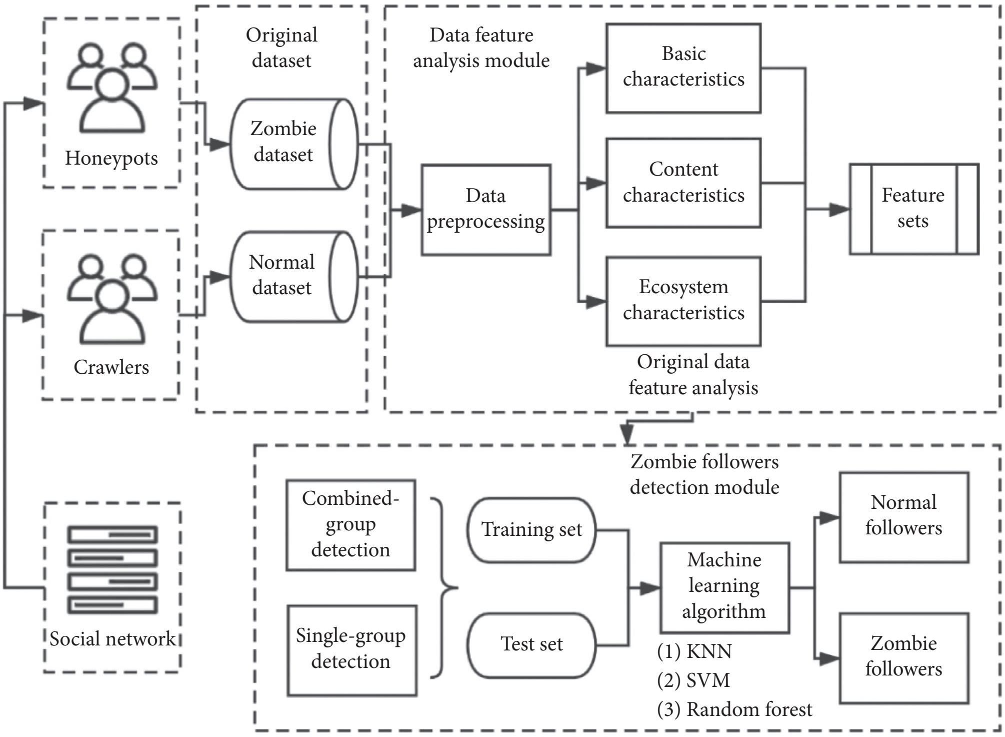

Figure 8: Structure of the detection model.

Table 10: Statistics of comparative experiment datasets.

| Collection | Total |
|---|---|
| 01_72****37 | 5063 |
| 02_72****04 | 943 |
| 03_72****91 | 6145 |
| 04_72****62 | 5257 |
| 05_72****21 | 5000 |
| 06_72****43 | 5022 |
| 07_72****47 | 5072 |
| 08_72****28 | 5105 |
| Normal users | 45559 |

features), although its detection accuracy was high, its recall, precision, and $F_1$ values were still generally low. In random forest detection, as an example, the accuracy of Groups 04 and 08 was 99.86% and 99.90%, respectively, but the recall, precision, and $F_1$ were 85.71%, 78.95%, and 82.19% for Group 04 and 99.58%, 99.37%, and 99.47% for Group 08. The large difference indicates that the richer the feature set is, the greater the promotion of recall, precision, and $F_1$ will be.

However, due to the feature differences among groups, the accuracy of the combined-group detection is generally lower than that of the single-group detection. It is worth noting that ecosystem features achieved approximately 90% accuracy in all single-group detections, but it decreased significantly in the combined-group detection. The random forest achieved the highest accuracy in both types of detection, with the combined-group detection achieving 98.75% accuracy (average accuracy of KNN and SVM is 99.4% and 99.46%, respectively). The table shows that Groups 03 and 06 with the longest registration time have the lowest accuracy. This indirectly confirms that the longer the zombie follower is cultivated, the more similar

Table 11: Detection accuracy of KNN for each feature category.

| KNN | Basic | Content | Ecosystem | All features | | | |
|---|---|---|---|---|---|---|---|
| | | | | Accuracy | Precision | Recall | $F_1$ |
| 01_72****37 | 98.55 | 92.05 | 90.16 | 99.57 | 96.48 | 99.30 | 97.87 |
| 02_72****04 | 99.83 | — | 97.89 | 99.82 | 95.90 | 95.41 | 95.65 |
| 03_72****91 | 99.07 | 91.36 | 87.27 | 99.14 | 95.04 | 97.66 | 96.34 |
| 04_72****62 | 99.77 | — | — | 99.77 | 54.84 | 70.83 | 61.82 |
| 05_72****21 | 99.89 | — | — | 99.89 | 50.00 | 100.00 | 66.67 |
| 06_72****43 | 98.00 | 93.25 | 67.06 | 97.99 | 88.88 | 90.60 | 89.73 |
| 07_72****47 | 99.28 | — | 80.28 | 99.36 | 94.72 | 99.00 | 96.81 |
| 08_72****28 | 98.74 | 90.59 | 90.74 | 99.64 | 97.11 | 99.16 | 98.12 |
| All data | 97.16 | 88.1792 | 66.32 | 97.59 | 95.56 | 97.85 | 96.69 |

Table 12: Detection accuracy of SVM for each feature category.

| SVM | Basic | Content | Ecosystem | All features | | | |
|---|---|---|---|---|---|---|---|
| | | | | Accuracy | Precision | Recall | $F_1$ |
| 01_72****37 | 98.38 | 92.32 | 90.21 | 99.55 | 97.01 | 98.48 | 97.74 |
| 02_72****04 | 99.82 | — | 97.96 | 99.83 | 96.76 | 94.71 | 95.72 |
| 03_72****91 | 99.25 | 98.08 | 88.10 | 99.06 | 98.71 | 93.32 | 95.94 |
| 04_72****62 | 99.86 | — | — | 99.86 | 78.95 | 85.71 | 82.19 |
| 05_72****21 | 99.93 | — | — | 99.93 | 73.68 | 93.33 | 82.35 |
| 06_72****43 | 98.19 | 92.64 | 90.26 | 98.28 | 94.30 | 87.59 | 90.82 |
| 07_72****47 | 99.38 | — | 91.62 | 99.43 | 95.33 | 98.99 | 97.12 |
| 08_72****28 | 98.89 | 92.11 | 90.84 | 99.73 | 98.52 | 98.63 | 98.57 |
| All data | 96.83 | 88.46 | 65.86 | 98.11 | 98.12 | 96.62 | 97.36 |

Table 13: Detection accuracy of Random forest for each feature category.

| Random forest | Basic | Content | Ecosystem | All features | | | |
|---|---|---|---|---|---|---|---|
| | | | | Accuracy | Precision | Recall | $F_1$ |
| 01_72****37 | 98.68 | 92.61 | 90.21 | 99.74 | 98.59 | 98.79 | 98.69 |
| 02_72****04 | 99.85 | — | 97.96 | 99.84 | 94.85 | 97.35 | 96.08 |
| 03_72****91 | 99.26 | 98.74 | 88.10 | 99.75 | 99.51 | 98.37 | 98.94 |
| 04_72****62 | 99.86 | — | — | 99.86 | 78.95 | 85.71 | 82.19 |
| 05_72****21 | 99.93 | — | — | 99.93 | 73.68 | 93.33 | 82.35 |
| 06_72****43 | 98.19 | 96.95 | 90.26 | 99.45 | 97.64 | 96.64 | 97.14 |
| 07_72****47 | 99.39 | — | 91.62 | 99.53 | 96.45 | 98.89 | 97.65 |
| 08_72****28 | 99.07 | 92.44 | 90.84 | 99.90 | 99.37 | 99.58 | 99.47 |
| All data | 97.16 | 92.11 | 65.86 | **99.14** | **99.03** | **98.57** | **98.80** |

it becomes to normal users and the more difficult its detection becomes.

## 6. Conclusion and Future Work

In this study, we focused on the features and detection of zombie followers from different companies. Through feature analysis, we described the current ecosystem of the zombie follower industry as follows: the ZF companies are constantly producing and cultivating zombie followers. Zombie followers that survive longer are more similar to normal users, thus lowering the detection rate by traditional methods. Furthermore, although the sources of zombie followers are different, the similar characteristics and the direct or indirect relationships between groups indicate that some zombie follower groups from different sources are actually controlled by the same organization. Finally, we used three different classification methods (i.e., KNN, SVM, and random forest) to detect zombie followers. The random forest performed the best with 99.14% accuracy. We also found that the richer the feature set, the greater the promotion of recall, precision, and $F_1$ of the detection results.

Interestingly, zombie follower services are only a small part of the black industry of malicious accounts. It also controls mass advanced social robot accounts, which have more advanced authentication and weight than normal users. Among them, some are compromised accounts and some have many real followers. In the future, we will conduct research on the features of such accounts.

## Data Availability

The data used to support the findings of this study are included within the manuscript and its supporting information files.

## Conflicts of Interest

The authors declare no conflicts of interest.

## Acknowledgments

This research was funded by the Sichuan Science and Technology Program, grant no. 2020YFG0076.

## Supplementary Materials

All relevant data are included within the manuscript and its supporting information files. (*Supplementary Materials*)

## References

[1] Sina IT, "Weibo posts the financial results for the fourth quarter and full year of 2019," 2020, https://tech.sina.com.cn/i/2020-02-26/doc-iimxyqvz6003265.shtml.

[2] S. Kumar, M. Jiang, T. Jung, R. J. Luo, and J. Leskovec, "MIS2: misinformation and misbehavior mining on the web," in *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, pp. 799-800, Los Angeles, CA, USA, February 2018.

[3] E. Ferrara, O. Varol, C. Davis, F. Menczer, and A. Flammini, "The rise of social bots," *Communications of the ACM*, vol. 59, no. 7, pp. 96–104, 2016.

[4] A. Thieltges, O. Papakyriakopoulos, J. M. Serrano, and H. Simon, "Effects of social bots in the Iran-debate on twitter," 2018, http://arxiv.org/abs/1805.10105.

[5] V. S. Subrahmanian, A. Azaria, S. Durst et al., "The DARPA twitter bot challenge," *Computer*, vol. 49, no. 6, pp. 38–46, 2016.

[6] E. E. Buckels, P. D. Trapnell, and D. L. Paulhus, "Trolls just want to have fun," *Personality and Individual Differences*, vol. 67, pp. 97–102, 2014.

[7] J. Cheng, M. Bernstein, C. Danescu-Niculescu-Mizil, and J. Leskovec, "Anyone can become a troll: causes of trolling behavior in online discussions," in *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*, pp. 1217–1230, Portland, OR, USA, February 2017.

[8] J. Cheng, C. Danescu-Niculescu-Mizil, and J. Leskovec, "Antisocial behavior in online discussion communities," in *Proceedings of the Ninth International AAAI Conference on Web and Social Media*, pp. 61–70, Oxford, UK, May 2015.

[9] S. Kumar, J. Cheng, J. Leskovec, and V. S. Subrahmanian, "An army of me: sockpuppets in online discussion communities," in *Proceedings of the 26th International Conference on World Wide Web WWW '17*, pp. 857–866, Perth, Australia, May 2017.

[10] S. K. Maity, A. Chakraborty, P. Goyal, and A. Mukherjee, "Detection of sockpuppets in social media," in *Proceedings of the Companion of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*, pp. 243–246, Portland, OR, USA, February 2017.

[11] M. Egele, G. Stringhini, C. Kruegel, and G. Vigna, "Towards detecting compromised accounts on social networks," *IEEE Transactions on Dependable and Secure Computing*, vol. 14, no. 4, pp. 447–460, 2017.

[12] E. Zangerle and G. Specht, "Sorry, I was hacked": a classification of compromised twitter accounts," in *Proceedings of the 29th Annual ACM Symposium on Applied Computing*, pp. 587–593, Gyeongju, South Korea, May 2014.

[13] J. P. Dickerson, V. Kagan, and V. S. Subrahmanian, "Using sentiment to detect bots on twitter: are humans more opinionated than bots?," in *Proceedings of the 2014 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, pp. 620–627, Beijing, China, August 2014.

[14] M. Jiang, P. Cui, A. Beutel, C. Faloutsos, and S. Yang, "Detecting suspicious following behavior in multimillion-node social networks," in *Proceedings of the 23rd International Conference on World Wide Web*, pp. 305-306, Seoul, South Korea, April 2014.

[15] H. Jiang, Y. Wang, and M. Zhu, "Discrimination of zombie fans on Weibo based on features extraction and business-driven analysis," in *Proceedings of the 17th International Conference on Electronic Commerce*, p. 13, Seoul, South Korea, August 2015.

[16] M. Fazil and M. Abulaish, "Identifying active, reactive, and inactive targets of socialbots in twitter," in *Proceedings of the International Conference on Web Intelligence*, pp. 573–580, Leipzig, Germany, August 2017.

[17] M. Jiang, P. Cui, and C. Faloutsos, "Suspicious behavior detection: current trends and future directions," *IEEE Intelligent Systems*, vol. 31, no. 1, pp. 31–39, 2016.

[18] S. Kumar and N. Shah, "False information on web and social media: a survey," 2018, http://arxiv.org/abs/1804.08559.

[19] S. Lehmann and P. Sapieżyński, "You're here because of a robot," 2020, http://sunelehmann.com/2013/12/04/youre-here-because- of- a-robot.

[20] S. Stieglitz, F. Brachten, D. Berthelé, M. Schlaus, C. Venetopoulou, and D. Veutgen, "Do social bots (still) act different to humans?–comparing metrics of social bots with those of humans," in *Proceedings of the International Conference on Social Computing and Social Media*, pp. 379–395, Vancouver, Canada, July 2017.

[21] Y. Ji, Y. He, X. Jiang, J. Cao, and Q. Li, "Combating the evasion mechanisms of social bots," *Computers & Security*, vol. 58, pp. 230–249, 2016.

[22] M. Khademi, S. Hosseini Moghaddam, and M. Abbaspour, "An empirical study of the effect of profile and behavioral characteristics on the infiltration rate of socialbots," in *Proceedings of the 2017 Iranian Conference on Electrical Engineering (ICEE)*, pp. 2200–2205, Tehran, Iran, May 2017.

[23] S. Lee and J. Kim, "WarningBird: a near real-time detection system for suspicious URLs in twitter stream," *IEEE Transactions on Dependable and Secure Computing*, vol. 10, no. 3, pp. 183–195, 2013.

[24] K. Thomas, C. Grier, J. Ma, V. Paxson, and D. Song, "Design and evaluation of a real-time URL spam filtering service," in *Proceedings of the 2011 IEEE Symposium on Security and Privacy*, pp. 447–462, Berkeley, CA, USA, May 2011.

[25] S. Santhosinidevi, “Towards detecting compromised accounts on social networks,” *International Journal for Research in Applied Science and Engineering Technology*, vol. 6, no. 4, pp. 71–73, 2018.

[26] Y. Shen, J. Yu, K. Dong, and K. Nan, “Automatic fake followers detection in Chinese micro-blogging system,” in *Advances in Knowledge Discovery and Data Mining*, pp. 596–607, Springer, Cham, Switzerland, 2014.

[27] C. Yang, R. Harkreader, J. Zhang, S. Shin, and G. Gu, “Analyzing spammers’ social networks for fun and profit: a case study of cyber criminal ecosystem on twitter,” in *Proceedings of the 21st International Conference on World Wide Web*, pp. 71–80, Lyon, France, April 2012.

[28] Y. Zhang and J. Lu, “Discover millions of fake followers in Weibo,” *Social Network Analysis and Mining*, vol. 6, no. 1, p. 16, 2016.

[29] C. S.-H. Eom, W. Lee, J. J.-H. Lee, and W.-S. Cho, “Find spammers by using graph structure,” in *Proceedings of the 2017 IEEE International Conference on Big Data and Smart Computing (BigComp)*, pp. 278-279, Jeju, South Korea, February 2017.

[30] J. Wang, X. He, Q. Gong, Y. Chen, T. Wang, and X. Wang, “Deep learning-based malicious account detection in the momo social network,” in *Proceedings of the 2018 27th International Conference on Computer Communication and Networks (ICCCN)*, Hangzhou, China, July 2018.

[31] H. Shen and X. Liu, “Detecting spammers on twitter based on content and social interaction,” in *Proceedings of the 2015 International Conference on Network and Information Systems for Computers*, pp. 413–417, Wuhan, China, January 2015.

[32] M. Ikram, L. Onwuzurike, S. Farooqi et al., “Measuring, characterizing, and detecting facebook like farms,” *ACM Transactions on Privacy and Security (TOPS)*, vol. 20, no. 4, p. 13, 2017.

[33] M. Héder, “A black market for upvotes and likes,” *Információs Társadalom*, vol. 19, no. 4, pp. 18–39, 2020.

[34] Cristofaro, E. De, A. Friedman, G. Jourjon, M. A. Kaafar, and M. Zubair Shafiq, “Paying for likes?: understanding facebook like fraud using honeypots,” in *Proceedings of the 2014 Conference on Internet Measurement Conference*, pp. 129–136, Vancouver, Canada, November 2014.

[35] K. Lee, C. James, and S. Webb, “Uncovering social spammers: social honeypots + machine learning,” in *Proceedings of the 33rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 435–442, Geneva, Switzerland, July 2010.

[36] X. Hu, J. Tang, and H. Liu, “Online social spammer detection,” in *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence AAAI’14*, pp. 59–65, Québec City, Canada, July 2014.

[37] K. Lee, B. David Eoff, and C. James, “Seven months with the devils: a long-term study of content polluters on twitter,” in *Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media*, Barcelona, Spain, July 2011.

[38] C.-H. Xia, H.-K. Li, and G.-Z. Sun, “Microblogging malicious user identification based on behavior characteristic analysis,” *Computer Science*, vol. 45, no. 12, pp. 111–116, 2018.

[39] S. Junyi, “Jieba: Chinese text segmentation,” 2020, https://github.com/fxsjy/jieba/.

[40] H. Li, “Analysis and implementation of spammers detection method based on social network,” in *CNKI*Beijing Jiaotong University, Beijing, China, 2020, http://cdmd.cnki.com.cn/Article/CDMD-10004-1017086622.htm.

[41] A. Andoni and P. Indyk, “Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions,” *Communications of the ACM*, vol. 51, no. 1, pp. 117–122, 2008.

[42] J. A. K. Suykens and J. Vandewalle, “Least squares support vector machine classifiers,” *Neural Processing Letters*, vol. 9, no. 3, pp. 293–300, 1999.

[43] L. Breiman, “Random forests,” *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001.