JOURNAL OF INFORMATION SCIENCE AND ENGINEERING 39, 489-506 (2023)
DOI: 10.6688/JISE.202305_39(3).0003

Characterizing Behaviors of Sockpuppets in Online Political Discourses*

HSIU-LING CHU[1], YU-CHEN DAI[1], WEI-BIN LEE[1,3], TING-SHEN HSIEH[2] AND MING-HUNG WANG[2,+]
[1]*Department of Information Engineering and Computer Science*
Feng Chia University
Taichung, 407 Taiwan
[2]*Department of Computer Science and Information Engineering*
National Chung Cheng University
Chiayi, 621 Taiwan
[3]*Information Security Research Center*
Hon Hai Research Institute, Hon Hai Precision Industry Co., Ltd.
Taipei, 114 Taiwan
E-mail: {purpleshah; double.dai.0129; donald880313}@gmail.com;
wei-bin.lee@foxconn.com; tonymhwang@cs.ccu.edu.tw

Recently, with the rapid development of social network services, political entities have employed social media to conduct political propaganda; they disseminate official announcements, promote candidates, and even attack the opposites. However, to influence public opinion, some of them will create sockpuppets to participate in discussions. To better understand the phenomenon, this study attempts to measure the behavioral differences between sockpuppets and normal users using network structure analysis and user behavior analysis. We observe and realize the behaviors of sockpuppets by collecting a dataset from PTT, a famous social forum in Taiwan. We propose three feature categories including author-based, commenter-based, and network-based features. From the analysis results, we find that sockpuppets are more active and attractive than ordinary users. Moreover, we observe that sockpuppets play an essential role in information flow from the results of the network structure of user relationships.

Keywords: information manipulation, political propaganda, social media, sockpuppets, social networks

1. INTRODUCTION

The popular and approachable social network services (SNS) have enabled political organizations to connect and interact with supporters and the general public. Online political propaganda has been conducted via SNS during election campaigns, urging citizens to participate in activities and vote [1]. For online users, SNS allows them to access news

Received January 8, 2022; revised March 2 & April 15, 2022; revised May 4, 2022.
Communicated by Po-Wen Chi.
+ Corresponding author.
* This work was supported by the Ministry of Science and Technology, Taiwan, under the Grant MOST 111-2222-E-194-003, MOST 111-2622-E-194-005, MOST 111-2927-I-194-001, MOST 110-2927-I-194-001, and MOST 107-2218-E-035-009-MY3.

of candidates, join discussions, and publish opinions easily. However, every coin has two sides. With the capabilities social networks have provided, threats that may undermine online democracy have emerged. For example, some entities may recruit organized users and develop programmable accounts to manipulate online opinions [2]. This can be achieved by publishing designated articles and flooding the platforms; thus, affecting people's judgment on political issues and election outcomes [3–5]. Beyond that, sockpuppets are an evolution of fake accounts. Sockpuppets are smarter than social bots [6], using human-like methods to hide in social networks to manipulate public opinion for a more favorable advantage. These tactics include forging verbal and feigning human behavior to misinterpret or fake posts, and responding quickly and enthusiastically.

To address this concern, in this study, we investigate the behavioral differences between sockpuppets and normal users utilizing network structure and user behavior analysis to identify bot-like programmable accounts and man-kind organized users. A dataset, including one-year-long online user activities in the most popular social forum in Taiwan, is collected and studied. The observation period spans over a national election in Taiwan. Another verified list of sockpuppets supported by the platform officials is also extracted and used to validate the sockpuppets activities of our discoveries. The major findings of this study are summarized as follows:

- **Publication, comments, and polarity of users.** From the user activities, we demonstrate the behavioral difference between normal users and verified sockpuppets. Significant differences are found in terms of publications, comments, and reactions of users.
- **Network structure of user relationship.** From the network structure, we apply multiple centrality algorithms to analyze sockpuppets. Consequently, we obtain that most sockpuppets greatly influence the social network.
- **Clustering phenomenon.** From the case observations after integrating features, we discern that the extent of clustering among sockpuppets appears to be higher than that among the normal users.

The remainder of the paper is organized as follows. In Section 2, we present related works about sockpuppet detection. In Section 3, we present our method in detail, including data collection and feature extraction. In Section 4, we demonstrate the analysis results of each feature. Finally, Section 5 presents the conclusion and discussion.

2. RELATED WORKS

With the rapid development of social networks, online political propaganda has been conducted and observed on social networking sites to approach Internet users. In 2011, Pariser *et al*. [7] discovered that fake news and selective media coverage caused a stratosphere phenomenon, making the polarization of ideology more serious. Additionally, Stukal *et al*. [8] pointed out that bot accounts on social platforms are related to disseminating malicious articles and fake news. These fake accounts controlled by malicious users for political propaganda can be divided into two categories according to whether they are programmable or manual: social bots and sockpuppets. Therefore, we will explore the differences between these two types of malicious users in the social bots and sockpuppets subsections.

2.1 Social Bots

Sevaral studies have shown that programmable accounts, also known as social bots, can influence public opinions [9, 10] and election campaigns [11, 12]. Thus, researchers have been developing methods to detect such social bots. They have also proposed several learning-based approaches. In 2016, Mehrotra *et al*. [13] used only graph-based centrality measures as the features of the learning model and achieved high accuracy on fake follower detection. In 2017, Cai *et al*. [14] proposed a deep learning detection model by integrating social behavior and content of accounts. Meanwhile, Zhang *et al*. [15] proposed that centrality is an essential indicator because it shows which node takes up a key position in one whole network. They presented degree, betweenness and closeness centralities from principle to algorithm. In 2018, Fazil *et al*. [16] used random forest, decision tree, and Bayesian networks to construct a high-performance spammer detection framework. In 2019, Shi *et al*. [17] combined the transition probability of clickstream sequences and semisupervised clustering to present a novel identification approach.

Great strides have been made by researchers in the study and detection of social bots. These detection methods have also made manipulators of such malicious users think about how to reduce the regularity in the programming, making it more like a normal account trying to evade and obfuscate all kinds of bots detection methods.

2.2 Sockpuppets

As the detection of social bots has become more accurate, the method for executing political propaganda by fake accounts has changed from automation to manual in recent years. In 2018, Wang *et al*. [6] observed that smarter sockpuppets use forged verbal and behavioral features to evade detection and management. In 2020, Schwartz *et al*. [18] combined machine learning and investigative journalism to examine sockpuppets from the inside. They interviewed a whistleblower that used to be a campaign member and exposed the activity of sockpuppets. In 2021, Pisciotta *et al*. [19] considered false identity accounts that look different but are created by the same person or organization. They manipulated these fake accounts to conduct inappropriate behaviors, such as manipulating opinions, spreading fake news, and disrupting other users. It is commonly known as the sockpuppet problem. Nguyen *et al*. [20, 21] found that time-series features have a significant effect on the recognition model. Additionally, Yu *et al*. [22] combined the association based on verbal and nonverbal features and proposed an adaptive multisource feature fusion method. Bandy *et al*. [23] propose a sockpuppet audit that characterize the effects of algorithmic curation on the source and topic diversities in Twitter timelines.

Various studies on sockpuppets have been fruitful. For example, Wang *et al*. [6] detected differences in sub-network graphs comparing community interactions and interest interactions between puppetmaster-sockpuppet and sockpuppet-ordinary; Schwartz *et al*. [18] detected sockpuppets by an internal news survey; Pisciotta *et al*. [19] focused on Internet science but must address false positives, and other research from the structure or algorithms of social media itself. Most of their research features belong to a single category, dominated by network features or text-language features.

From the above research perspective, referred to as sockpuppet accounts, it has the following characteristics: creating multiple accounts for the same user or organization, using these accounts to publish articles or comments to influence Internet public opinion,

forging words and behavior, and spreading fake news to mislead the public.

However, studies about detecting manually operated accounts remain scant because their behaviors may not be as regular as programmable ones. In this paper, we analyze the differences between sockpuppets and normal users from three perspectives: author, commenter, and network structure. We also attempt to describe the activity of sockpuppets and propose useful features for identification.

3. METHODS

In this section, we introduce the proposed framework for sockpuppet analysis. The framework includes data collection, feature extraction, and feature analysis. In this work, we take the sockpuppet identification task as a binary classification problem where every user can be classified into sockpuppets and ordinary users.

3.1 Data Collection

We collect our dataset from one of the most predominant social platforms in Taiwan, the "Gossiping" board of the PTT Bulletin Board System, to investigate the sockpuppet accounts and potential information manipulation behavior during elections. For the Gossiping board, in the history board of PTT, the ranking of the most popular boards will be announced yearly. We have excerpted the top 5 from the latest announcement of [Site History] 2021 Most Popular Rankings of PTT Boards[1], as presented in Table 1, where the numbers below the year represent the highest influx of people on the board at a point in time.

Table 1. 2021 most popular rankings of PTT boards.

RK	△▽	BOARD	2017	2018	2019	2020	2021
1	–	Gossiping	40,464	106,728	39,028	65,370	66,628
2	New	Olympics_ISG	13,280	7,070	104	–	47,152
3	△2	Stock	6,192	12,232	6,336	17,898	27,359
4	–	NBA	35,933	35,797	37,957	20,197	23,061
5	▽3	Golden-Award	21,784	24,853	23,948	25,154	21,655

The top-ranking is the Gossiping board, the highest popularity figure in 2018 was 106,728, which occurred at 2:32 on November 25, and the event was the final vote by the mayor of Taipei. The ecology of PTT, Taiwanese political comments are mostly gathered on Gossiping board, therefore, we chose the Gossiping board as the source of data collection. During the one-year-long observation, from July 1, 2018, to July 1, 2019, the dataset collection spans over the 2018 local election in Taiwan. In the dataset, every entry is consists of an article and its comments and metadata. The detailed information consists of the following items:

(i) **Article author information:** author ID, screen name, and IP address;

(ii) **Article metadata:** article ID, and post time;

[1] https://www.ptt.cc/bbs/PttHistory/M.1641044847.A.4A7.html

(iii) **Article content:** the text content of the article;

(iv) **Commenter's comment and rating:** comment body, comment time, and a positive/neutral/negative rating accompanied by the comment.

However, to retrieve a creditable source of sockpuppets accounts, we collected articles from another board, "ID_Multi," where users report suspicious accounts, and the official announces judgment. We manually read and extract the sockpuppet accounts according to the official announcement and use it as the ground truth of the paper.

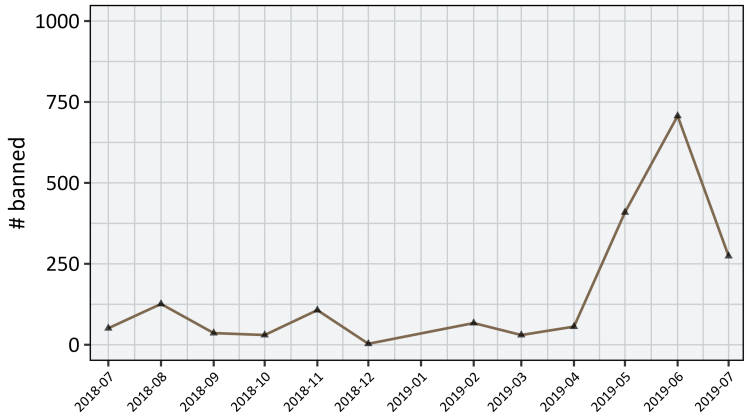

Fig. 1. The monthly number of accounts being banned during observation.

Fig. 1 shows the monthly number of banned users from 2018 to 2019 and a PTT for sockpuppets. In August 2018, Typhoon Jebi struck Japan, resulting in the closure of Kansai Airport. During the period, a Taiwanese diplomat was criticized for handling requests from Taiwanese. However, a few days later, the diplomat committed suicide[2]. After official investigations, evidence shows that some criticisms may have been manipulated by accounts on PTT. PTT officials conducted a large-scale operation to ban accounts in November 2019, and about 6,000 users were declared sockpuppets. Therefore, we believe that the mechanism of "ID_Multi" banning accounts can provide a sufficient basis of trust.

Fig. 2 shows the framework diagram of our research method.

We divide the research section into three main stages: data collection, data processing, and feature analysis. In the data collection stage, we collect data from ID_Multi and Gossiping boards in the PTT bulletin board system. In the data processing stage, the normal and sockpuppet user types are manually labeled from the ID_Multi board. Meanwhile, the author-based, commenter-based, and network-based features are extracted from the Gossiping board. Then, they are integrated into our user-feature-label data frame. In the final stage, we analyze the features of network structure and user behavior.

[2]`https://w.wiki/ara`

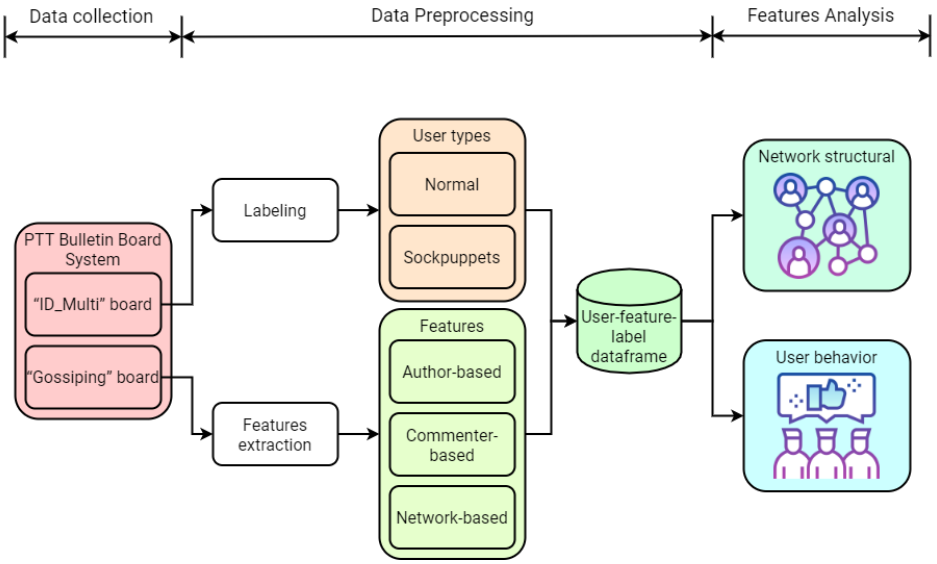

Fig. 2. Research methodology framework.

Table 2 summarizes of our dataset.

Table 2. A summary of our dataset.

	author		commenter			
Type	# account	# article	# account	# comment	# like	# boo
Normal	41,454	737,000	223,280	29,131,228	13,226,884	4,323,511
Sockpuppet	1,357	62,393	2,469	742,693	352,553	116,312

3.2 Features Extraction

We extract article attributes from the data collection and develop corresponding features to indicate each user's publishing and commenting behaviors. These features are created from author-based, commenter-based, and network-based perspectives. Table 3 presents the feature list, feature categories, and details.

1) Author-based features: From the previous studies, sockpuppets are used to publish for shaping the opinion. They endeavor to produce or echo particular voices to influence other users. We introduce a set of features in this section to describe author activity behaviors.

- *The number of pushes received (A1), the number of boos received (A2), and the number of neutrals received (A3):* There are three emotion ratings on PTT. When a user comments on the article, users can choose one emotion rating to express their emotional attitude. Emotion rating includes "positive (#push)," "negative (#boo)," and "neutral." For example, users can comment with #push to express their agreement with the content of the article. Additionally, a "neutral" rating enables users to choose a neutral attitude in the article. These three features can be used to evaluate the public opinion toward the user's articles. For each article a by

Table 3. Author, commenter, and network features.

Features	Categories	Description
A1	Author-based	The number of pushes received
A2	Author-based	The number of boos received
A3	Author-based	The number of neutrals received
A4	Author-based	Author polarity score
A5	Author-based	The number of "Exposing" posted
A6	Author-based	The number of "Gossiping" posted
A7	Author-based	The number of "News" posted
A8	Author-based	The longest consecutive posting days
A9	Author-based	The average number of article posted on weekday per day
A10	Author-based	The average number of article posted on weekend per day
C1	Commenter-based	The number of pushes given
C2	Commenter-based	The number of boos given
C3	Commenter-based	The number of neutrals given
C4	Commenter-based	Commenter preference score
C5	Commenter-based	The Median of response time
C6	Commenter-based	The longest consecutive commenting days
C7	Commenter-based	The average number of comments given on weekday per day
C8	Commenter-based	The average number of comments given on weekend per day
N1	Network-based	In-degree centrality
N2	Network-based	Out-degree centrality
N3	Network-based	In-closeness centrality
N4	Network-based	Out-closeness centrality
N5	Network-based	Betweenness centrality
N6	Network-based	PageRank centrality
N7	Network-based	Eigenvector centrality

author u, three ratings received of author u is denoted as $push.r_u = \sum_{a \in u} push_a$, $boo.r_u = \sum_{a \in u} boo_a$, and $neutral.r_u = \sum_{a \in u} neutral_a$. Where $push_a$, boo_a, and $neutral_a$ are the numbers of positive, negative, and neutral ratings given by the commenters, respectively.

- *Author's polarity score (A4):* This feature describes how online users react to an author's publications. This metric attempts to quantify the online audience's attitude toward the author. The polarity score of author u is derived by subtracting $boo.r_u$ from $push.r_u$, as defined in Eq. (1).

$$Polarity_u = push.r_u - boo.r_u \tag{1}$$

- *The number of "Exposing" posted (A5), the number of "Gossiping" posted (A6), and the number of "News" posted (A7):* In PTT, each board has its rules to stipulate for the article categories, such as there are five article categories on the "Gossiping" board. When users want to publish the article on the "Gossiping" board, they must

choose one article category to express the intention of the article. Article categories include "Exposing," "Gossiping," "News," "Wanted/missing," and "Public Service Announcement (PSA)." For example, users can post with #exposing to mean that they want to expose something. They can post with #gossiping to show that they want to gossip about some topic. Additionally, the "news" category enables users to share the news from local media. These three features can be used to evaluate the intention of the user's article publication. Since "Wanted/missing" and "PSA" are used less and usually used by the officials, we do not use them.

- *The longest of consecutive posting days (A8):* This feature is used to describe the behavior of users who post articles consecutively. According to the post time, we sort each user's articles and count the longest consecutive post days for each user.
- *The average number of articles posted on weekdays per day (A9) and the average number of articles posted on weekends per day (A10):* These two features are used to measure whether users regard article publishing as daily work. For each article a by author u, we denote the number of posts on weekdays and weekends as $daily.avg_{a.weekday}$ and $daily.avg_{a.weekend}$. For each author u, the total number of activity days on weekdays and weekends are defined as $sum.weekdays_u$ and $sum.weekends_u$, respectively. Features A9 and A10 are calculated as Eqs. (2) and (3), respectively.

$$A9_u = \frac{\sum_{a \in u} daily.avg_{a.weekday}}{sum.weekdays_u} \tag{2}$$

$$A10_u = \frac{\sum_{a \in u} daily.avg_{a.weekend}}{sum.weekends_u} \tag{3}$$

2) Commenter-based features: Sockpuppets can also behave as commenters to express opinions by replying; this section proposes a list of features to describe the behavior when users act as commenters.

- *The number of pushes given (C1), the number of boos given (C2), and the number of neutrals given (C3):* These three features are similar to the number of three ratings received (A1, A2, and A3) that we have denoted. We turn our view from an author to a commenter, summing up each user's commenter's emotion ratings. We can use these three features to evaluate the individual opinion of the commenter. For each comment r by commenter c, three ratings given by commenter c are denoted as $push.g_c = \sum_{r \in c} push_r$, $boo.g_c = \sum_{r \in c} boo_r$, and $neutral.g_c = \sum_{r \in c} neutral_r$. Where $push_r$, boo_r, and $neutral_r$ are the number of positive, negative, and neutral ratings given by the commenters, respectively.
- *Commenter preference score (C4):* We define the commenter preference score to measure the attitude from a user to other articles. The preference score of commenter c is derived by subtracting $boo.g_c$ from $push.g_c$, as defined in Eq. (4).

$$Preference_c = push.g_c - boo.g_c \tag{4}$$

- *The median of response time (C5):* We use the median of response time to describe whether the user focuses on responding to the article. For each comment r to article p, we denote the time of the comment as $time_r$, and the time of article publishing as $time_p$. The response time is defined as the user's comment time minus the article post time, as shown in Eq. (5).

$$Response.time = time_r - time_p \tag{5}$$

- *The longest consecutive commenting day (C6):* This feature describes the behavior of the user who consecutively comments on the article. We sort each user's comment by a day and count the longest consecutive comment days.

- *The average number of comments given on weekday per day (C7) and the average number of comments given on weekend per day (C8):* These two features are similar to features A9 and A10. We turn our view from an author to a commenter and calculate the average number of comments given on weekdays and weekends per day.

3) Network-based features: From the previous studies, using graph-based centrality measures as features can yield high model performance on fake follower detection [13]. Each centrality has its definition. The difference is that the previous study applies these centralities to analyze the relationship between specified followed accounts and their follower accounts. Meanwhile, we take articles as research subjects and focus on the relationship between authors and commenters more nuanced interactions between them. In this section, we introduce a set of features to describe network structure. Additionally, we set each comment as edge E and each user as vertex V and then create a directed graph $G(V,E)$. The direction between each vertex is from commenter to author.

- *In-degree centrality (N1) and out-degree centrality (N2):* In-degree centrality is defined as the number of edges E coming into a vertex V in graph G. Out-degree centrality is defined as the number of edges E coming out from a vertex V in graph G. Moreover, we normalize the features as divided by $n-1$, where n is the number of vertex V in the graph G.

- *In-closeness centrality (N3) and out-closeness centrality (N4):* The closeness centrality of a vertex V can be defined as the reciprocal of the average length of the shortest paths to/from all the other vertices in the graph G. We normalize the features by multiplying the raw closeness by $N-1$, where N is the number of vertex V in the graph G, as shown in Eq. (6).

$$Closeness.c(V_i) = \frac{N-1}{\sum_{V_j} dist(V_i,V_j), V_i != V_j}, \tag{6}$$

  where $dist(V_i,V_j)$ is the distance between vertices V_i and V_j.

- *Betweenness centrality (N5):* The betweenness centrality of a vertex V_i is the number of the all-pairs shortest path passing through V_i, as shown in Eq. (7).

$$Betweenness.c(V_i) = \frac{dist.n_{x,y}(V_i)}{dist.n_{x,y}}, \tag{7}$$

where $dist.n_{x,y}$ is the total number of shortest paths from vertex x to vertex y and $dist.n_{x,y}(V_i)$ is the number of the paths which pass through V_i.

- *PageRank centrality (N6):* PageRank is a classic algorithm for ranking web pages [24]. It refers to the possibility of a web page being browsed. Each web page has its PageRank, depending on the link relationship between web pages so that popular web pages stand out from the crowd. In this study, we adopted PageRank to measure the user's influence. The PageRank of vertex V_i is computed using Eq. (8).

$$PR.c(V_i) = d \sum_{V_j \in D.in(V_i)} \frac{PR.c(V_i)}{D.out(V_j)} + \frac{1-d}{N}, \tag{8}$$

where N is the total number of vertex V and d is the damping factor, which is commonly set to 0.85 [24]. $D.in(V_i)$ is a set of vertex, which link to V_i, and $D.out(V_j)$ is the out-degree of V_j.

- *Eigenvector centrality (N7):* We define the interaction-based eigenvector centrality to measure the relationship between sockpuppet and normal account. Our method is based on power iteration [25], as shown in Eqs. (9)-(11).

$$Eigenvector.c_n^{(k)} = A \times \lambda, \tag{9}$$

$$Eigenv.norm_n^{(k)} = \frac{Eigenvector.c_n^{(k)}}{\sqrt{\sum_{i=1}^{n} (Eigenvector.c_i)^2}}, \tag{10}$$

$$Eigenvector.c_n^{(k+1)} = A \times Eigenv.norm_n^{(k)}, \tag{11}$$

where A is an adjacency matrix created by graph G, λ is an $n \times 1$ matrix created by the label of the user such that its element λ_n is 1 when the user is verified as a sockpuppet, and 0 when the user is legitimate. Moreover, n is the total number of accounts, and k represents the number of iterations starting from 1.

4. ANALYSIS

4.1 Sockpuppet Announcements

As described in Subsection 3.1 and Fig. 1, in addition to the Typhoon Jebi in Japan, there were elections for local public officials in Taiwan in 2018. We observe the characteristics related to sockpuppets from the ban mechanism of the ID_Multi board, such as pretending to be a general user with forged words and behaviors, driving ethos and conducting public opinion manipulation. In this study, we try to figure out how different the behavior of sockpuppets is from that of normal users during sockpuppets manipulated information.

4.2 Features Analysis

We discuss in this subsection features proposed in Subsection 3.2. To present the major differences between sockpuppets and normal users, we use the empirical cumulative distribution function (ECDF). Then, we summarize several classifications of behavior as follows:

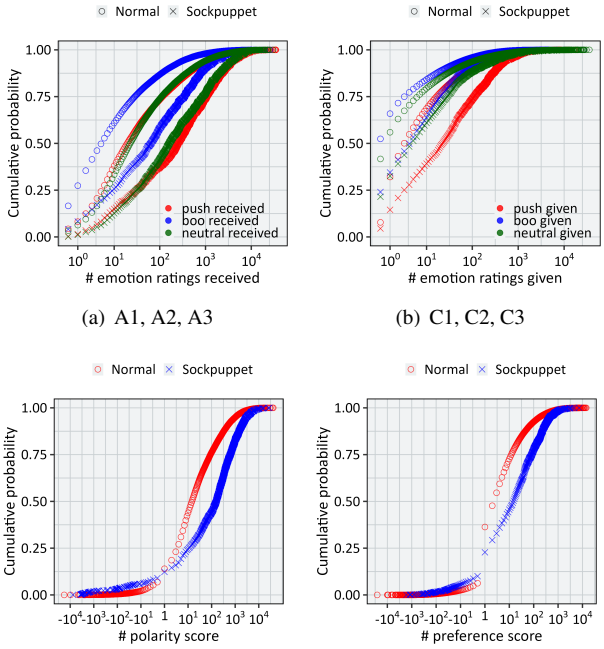

(a) A1, A2, A3 (b) C1, C2, C3

(c) A4 (d) C4

Fig. 3. The empirical cumulative distribution function (ECDF) for sockpuppets and normal users describe by rating-based behavior.

(i) **Rating-based behavior:** Fig. 3 (a) shows the normalized cumulative results of the number of received and (b) of the given, three emotions: push, boo, and neutral. We illustrate with the cumulative number of pushes received by 50% of the sockpuppet and the normal user in Fig. 3 (a), the sockpuppet is significantly higher than the normal user by an order of magnitude.

Figs. 3 (a) and (b) show that sockpuppets are greater than normal users in terms of the number of each emotion rating received or given. In Fig. 3 (b), about 50%-60% of sockpuppets received more than 100 positive/negative/neutral ratings from the public. Nevertheless, this proportion is about 20%-25% for normal users, which is significantly lower than sockpuppets. Fig. 3 (b) also shows that sockpuppets prefer to give positive comments. More than 25% of sockpuppet accounts give more than 100 positive ratings; nevertheless, this proportion is only 10% for normal users. This result indicates that sockpuppets' articles have received more attention and public reactions then normal users. Sockpuppets are also glad to comment on articles, especially using push (positive) emotion rating.

Figs. 3 (c) and (d) show the results of polarity and preference scores. We can observe that sockpuppets are higher than normal users because of positive and negative values, indicating that sockpuppets are more polarized and can shape opinions to attract discussions.

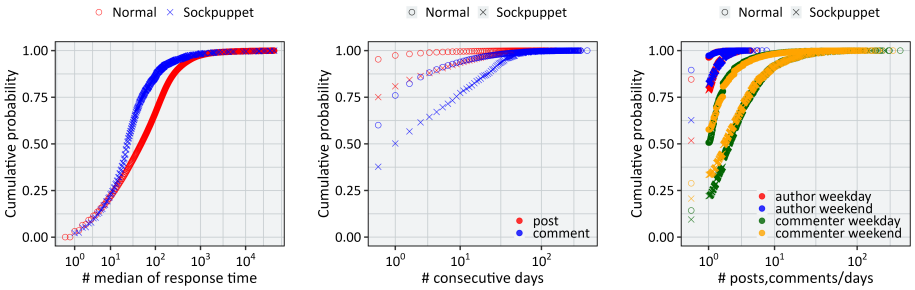

(a) C5 (b) A8, C6 (c) A9, A10, C7, C8

Fig. 4. The empirical cumulative distribution function (ECDF) for sockpuppets and normal users describe by time-based behavior.

(ii) **Time-based behavior:** We analyze in this section whether the sockpuppets regard publishing or commenting on the articles as a daily duty and a job. Fig. 4 (a) compares the response times to posts between sockpuppets and normal users, and shows that sockpuppets' response time is faster than that of normal users. About 62.5% of sockpuppets respond to the articles in less than an hour; however, this proportion is only 37.5% for normal users. Fig. 4 (b) shows whether users publish or comment on articles daily. We can observe that sockpuppets' consecutive activity days are longer than those of normal users. The above analysis explains that sockpuppets focus more on the discussion in the forum and comment rapidly. Additionally, we assume that if sockpuppets regard publishing or commenting on articles as a job, they may have fixed operating hours. However, we cannot find that behavior we assume from Fig. 4 (c), because from a comparison of the number of daily posts and comments on weekdays and weekends, there is not enough evidence to support the hypothesis that sockpuppets are recruited.

(iii) **Network structure:** From the directionality of vertices and edges in the directed graph, we can understand the interactive relationship between accounts and observe the trend of each category in several centralities. Fig. 5 (a) shows the probability distribution of in-degree centrality, representing the situation where the account has been connected by others. We observe that about 80% of the normal users are 0 connected, and about 55% of the sockpuppet accounts are connected. Fig. 5 (b) shows the out-degree centrality situation, indicating that the account is connected to others. About 76% of normal users are in 0 connection status, and the status of sockpuppets is similar to N1. Fig. 5 (c) shows the overall connection of the account. In this case, about 65% of normal users do not have any network interaction, whereas only about 40% of sockpuppets are in a similar situation. As shown in Fig. 5, we can infer that sockpuppets are active groups in social networks.

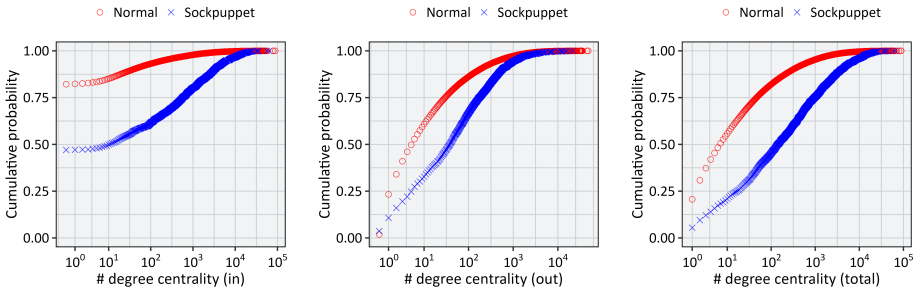

(a) N1 (b) N2 (c) Total degree

Fig. 5. The empirical cumulative distribution function (ECDF) for sockpuppets and normal users describe by network-based features-I.

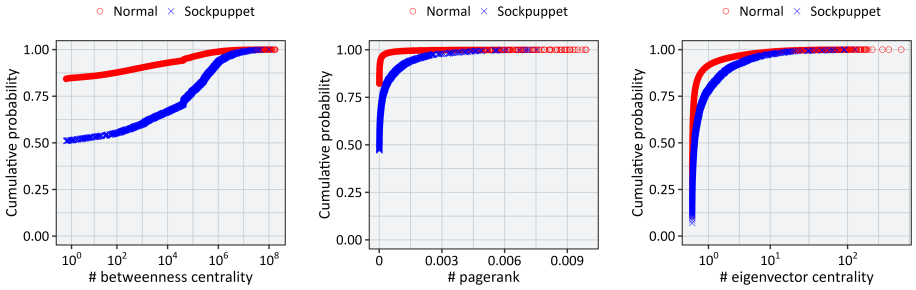

(a) N5 (b) N6 (c) N7

Fig. 6. The empirical cumulative distribution function (ECDF) for sockpuppets and normal users describe by network-based features-II.

Fig. 6 (a) shows the distribution of betweenness centrality by two types of users. The figure also shows that the betweenness centrality of sockpuppets is higher than that of normal users. Fig. 6 (b) shows the PageRank centrality distribution by two types of users. The figure also shows that less than 50% of sockpuppets' PageRank centrality equal 0 compared with that of normal users. However, about 80% of normal users' PageRank centrality equal 0. These results indicate that sockpuppets play an essential role in the social network structure. They may be the opinion leaders in the community or the key part of information flow. Furthermore, Fig. 6 (c) shows the distribution of our proposed interaction-based eigenvector centrality. As shown in the figure, sockpuppets and normal users are distinguished, indicating that undetected sockpuppets can be identified using the interaction of verification sockpuppets and suspicious users.

Additionally, we randomly selected five sockpuppets from the sample of sockpuppets and drew the network connection diagram extended from them (Fig. 7). Here orange and blue represent sockpuppets and normal users, respectively. It can be seen from the size and number of the dots that a few sockpuppets have a strong influence on a large number of normal users. This may be similar to the emotion rating analysis in Fig. 3 (b), which is more popular with the masses.

Fig. 8 shows our conjecture for sockpuppet and general users from the perspective of ego networks. Before selecting cases, we filtered out users with positive emotion values (push) less than 10 to reduce unsuitable (such as static accounts) usage. We randomly selected a target account from each of 292 sockspuppets and 71,099 normal users, and plotted them with order = 1, as shown in the figure. In the ego network of the targeting sockpuppet, we observed that six other sockpuppets are connecting with the target sockpuppet, accounting for 2.4% of all sockpuppets. On the contrary, the ratio of normal users connecting with the target sockpuppet is less than 0.0004%. These statistics imply that sockpuppets are more active with other sockpuppets rather than normal users.

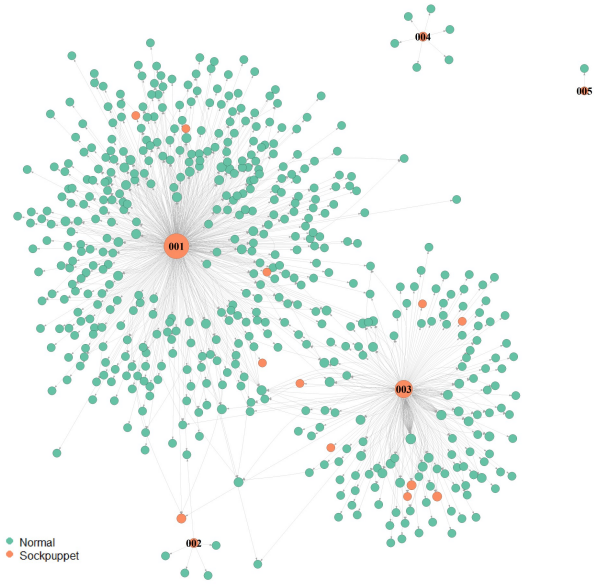

Fig. 7. Network connection example based on sockpuppets.

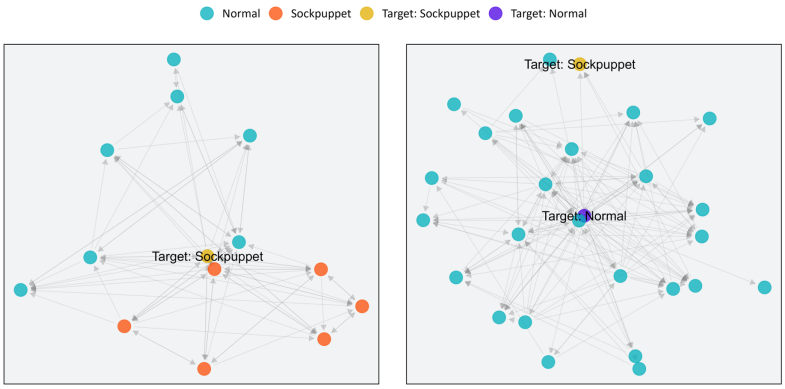

Fig. 8. Network connection example based on ego network.

5. CONCLUSION

Information manipulation and political propaganda have become crucial issues with the emergence of social networks. In this paper, we analyze the user behavior in a one-year-long observation on a popular social forum in Taiwan. We extract three kinds of features from the perspectives of author, commenter, and network perspectives. The major findings are summarized as follows:

- **Sockpuppets are more active and popular.** The articles published by sockpuppets are more concerned by other users than those published by normal users, affecting public opinion. We also observe that sockpuppets are more active than normal users, whether they are the author's identity or of commenters.
- **Most sockpuppets play an essential role in information flow.** We find that each centrality score of sockpuppets is higher than that of normal users through network structure analysis. The results show that most sockpuppets play a tremendous part in information flow and probably are key opinion leaders in the community.
- **Sockpuppets have a crisis of manipulating public opinion and influencing society.** This will lead to public opinion being led to a specific purpose and disseminated, benefiting a certain group or individual while also undermining freedom of speech. Ultimately, it will only get worse for people to get lost in the torrent of public opinion manipulation.

From Subsection 2.2 we find that most research features on sockspuppets are of a single category, dominated by network or text-language features. The special features of this study is that our distinguishing features fuse author-based, commenter-based, and network-based. From the integrated study of emotional value, active time, centrality interaction, and page ranking, we can quickly detect the presence of sockpuppets and their corresponding clusters from the features. Based on this study, we provide a broader approach in feature extraction for political research and social sciences.

In this paper, we familiarize ourselves with data science through statistical analysis. In future works, we will explore the sockpuppet relationship from a network perspective and the similarities between sockpuppets' content and that of normal users, such as the speculation about the ego network in Subsection 4.2, we will move towards discovering sockpuppet interaction and clustering patterns, and develop detection mechanisms that can be applied to various group networks. Additionally, we will also propose a classifier for identifying sockpuppets. For example, we will try to develop a sockpuppet account classifier using a graph representation learning algorithm. We will use machine learning and data science methods to maximize the effect of the proposed features. We hope the outcomes of this study can help improve the transparency of discussions, uncover the manipulation groups, and retain the online democracy.

REFERENCES

1. S. Bradshaw and P. N. Howard, "Challenging truth and trust: A global inventory of organized social media manipulation," *The Computational Propaganda Project*,

Vol. 1, 2018.
2. A. Bessi and E. Ferrara, "Social bots distort the 2016 US presidential election online discussion," *First Monday*, Vol. 21, 2016.
3. G. Pennycook, T. D. Cannon, and D. G. Rand, "Prior exposure increases perceived accuracy of fake news," *Journal of Experimental Psychology: General*, Vol. 147, 2018, p. 1865.
4. S. Aral and D. Eckles, "Protecting elections from social media manipulation," *Science*, Vol. 365, 2019, pp. 858-861.
5. A. Deb, L. Luceri, A. Badaway, and E. Ferrara, "Perils and challenges of social media and election manipulation analysis: The 2018 US midterms," in *Companion Proceedings of World Wide Web Conference*, 2019, pp. 237-247.
6. J. Wang, W. Zhou, J. Li, Z. Yan, J. Han, and S. Hu, "An online sockpuppet detection method based on subgraph similarity matching," in *Proceedings of IEEE International Conference on Parallel & Distributed Processing with Applications, Ubiquitous Computing & Communications, Big Data & Cloud Computing, Social Computing & Networking, Sustainable Computing & Communications*, 2018, pp. 391-398.
7. E. Pariser, *The Filter Bubble: What the Internet is Hiding From You*, Penguin, UK, 2011.
8. D. Stukal, S. Sanovich, J. A. Tucker, and R. Bonneau, "For whom the bot tolls: a neural networks approach to measuring political orientation of twitter bots in russia," *Sage Open*, Vol. 9, 2019, No. 2158244019827715.
9. L. Hagen, S. Neely, T. E. Keller, R. Scharf, and F. E. Vasquez, "Rise of the machines? examining the influence of social bots on a political discussion network," *Social Science Computer Review*, Vol. 40, 2022, pp. 264-287.
10. P. N. Howard, *Lie Machines: How to Save Democracy from Troll Armies, Deceitful Robots, Junk News Operations, and Political Operatives*, Yale University Press, CT, 2020.
11. T. R. Keller and U. Klinger, "Social bots in election campaigns: Theoretical, empirical, and methodological implications," *Political Communication*, Vol. 36, 2019, pp. 171-189.
12. K.-C. Yang, P.-M. Hui, and F. Menczer, "Bot electioneering volume: Visualizing social bot activity during elections," in *Companion Proceedings of World Wide Web Conference*, 2019, pp. 214-217.
13. A. Mehrotra, M. Sarreddy, and S. Singh, "Detection of fake twitter followers using graph centrality measures," in *Proceedings of the 2nd International Conference on Contemporary Computing and Informatics*, 2016, pp. 499-504.
14. C. Cai, L. Li, and D. Zengi, "Behavior enhanced deep bot detection in social media," in *Proceedings of IEEE International Conference on Intelligence and Security Informatics*, 2017, pp. 128-130.
15. J. Zhang and Y. Luo, "Degree centrality, betweenness centrality, and closeness centrality in social network," in *Proceedings of the 2nd International Conference on Modelling, Simulation and Applied Mathematics*, Vol. 132, 2017, pp. 300-303.
16. M. Fazil and M. Abulaish, "A hybrid approach for detecting automated spammers in twitter," *IEEE Transactions on Information Forensics and Security*, Vol. 13, 2018, pp. 2707-2719.

17. P. Shi, Z. Zhang, and K.-K. R. Choo, "Detecting malicious social bots based on clickstream sequences," *IEEE Access*, Vol. 7, 2019, pp. 28855-28862.
18. C. Schwartz and R. Overdorf, "Disinformation from the inside: Combining machine learning and journalism to investigate sockpuppet campaigns," in *Companion Proceedings of the Web Conference*, 2020, pp. 623-628.
19. G. Pisciotta, M. Somenzi, E. Barisani, and G. Rossetti, "Sockpuppet detection: a telegram case study," *arXiv Preprint*, 2021, arXiv:2105.10799.
20. N.-L. Nguyen, M.-H. Wang, Y.-C. Dai, and C.-R. Dow, "Understanding malicious accounts in online political discussions: A multilayer network approach," *Sensors*, Vol. 21, 2021, p. 2183.
21. N.-L. Nguyen, M.-H. Wang, and C.-R. Dow, "Learning to recognize sockpuppets in online political discussions," *IEEE Systems Journal*, Vol. 16, 2022, pp. 1873-1884.
22. H. Yu, F. Hu, L. Liu, Z. Li, X. Li, and Z. Lin, "Sockpuppet detection in social network based on adaptive multi-source features," *Modern Industrial IoT, Big Data and Supply Chain*, 2021, pp. 187-194.
23. J. Bandy and N. Diakopoulos, "More accounts, fewer links: How algorithmic curation impacts media exposure in twitter timelines," *Proceedings of ACM on Human-Computer Interaction*, Vol. 5, 2021, pp. 1-28.
24. S. Brin and L. Page, "The anatomy of a large-scale hypertextual web search engine," *Computer Networks and ISDN Systems*, Vol. 30, 1998, pp. 107-117.
25. N. Meghanathan, "A comprehensive analysis of the correlation between maximal clique size and centrality metrics for complex network graphs," *Egyptian Informatics Journal*, Vol. 22, 2021, pp. 339-355.

Hsiu-Ling Chu received her MS degree in Information Engineering at Feng Chia University in 2010. She is currently pursuing her Ph.D. degree in the Department of Information Engineering and Computer Science, Feng Chia University, Taichung City, Taiwan. Her research interests include data science, deep learning, and information security.

Yu-Chen Dai is an Engineer at the National Cybersecurity Center of Excellence, Taiwan. He received his master's degree in Information Engineering at Feng Chia University in 2021. His research interests are the analysis of malicious behavior on social media, including information manipulation and disinformation. Currently, he focuses on using network structure and user behavior to identify the sockpuppets in social media.

Wei-Bin Lee received his Ph.D. degree at National Chung Cheng University in 1997. Dr. Lee served as a Professor in the Department of Information Engineering and Computer Science at Feng Chia University. He was also a Visiting Professor at both Carnegie Mellon University in the USA and University of British Columbia in Canada. Since 2021, he has served as the CEO of HonHai Research Institute as well as the Director of the Information Security Research Center. Before joining the Foxconn Group, the CEO Wei-Bin Lee held important positions in Taipei City Government, Taipei Fubon Bank, Fubon Financial Holdings, and Feng Chia University. He has a deep research foundation in network security, cryptography, digital rights management, and privacy/security management and governance, as well as practical experience.

Ting-Shen Hsieh received the BS degree in Information Engineering at National Taiwan Ocean University in 2021. Currently, he is studying for the MS degree in the Department of Information Engineering and Computer Science, National Chung Cheng University, Chiayi, Taiwan.

Ming-Hung Wang is an Assistant Professor in the Department of Computer Science and Information Engineering at National Chung Cheng University. He received his Ph.D. in Electrical Engineering from National Taiwan University in 2017. He leads the Digital Society and Security Lab (DIGI-SSL), focusing on issues related to computational social science and information security.