





#WhoDefinesDemocracy: Analysis on a 2021 Chinese Messaging Campaign

Charity S. Jacobs^(✉)  and Kathleen M. Carley 

CASOS, Institute for Software Research Carnegie Mellon University,
Pittsburgh, PA 15213, USA
{csking, carley}@andrew.cmu.edu

Abstract. China has embraced the social media domain to promote pro-Chinese narratives and stories in recent years. However, China has increasingly been accused of launching information operations using methods such as bot activity, puppet accounts and other forms of inauthentic activity to amplify pro-Chinese messaging. This paper provides a comprehensive network analysis characterization of the hashtag influence campaign China promoted against the US-hosted Summit on Democracy in December 2021, in addition to methods to identify different types of actors within this type of influence campaign.

Keywords: Chinese influence operations · Bots · Network Analysis

1 Introduction

Social media is a critical domain for connecting and promoting ideas and discussion at the international level. China is a relative newcomer to the world of public relations and messaging to the international community, passing domestic legislature as recently as 2007 to increase the country's discourse power by creating a Chinese-controlled media platform Xinhua, capable of shaping narratives about China at the global level [3]. Increasingly, Chinese Communist Party (CCP) officials are using western social media platforms to defend Chinese national interests. These official accounts are not by themselves unique from any other country's use of social media, but have been used to spread disinformation and conspiracy theories such as the origins of COVID-19 [7] and human rights abuse in Xinjiang [8].

Recent research has shown China utilizes Twitter messaging to target western audiences and generally refrains from the use of memes that often take on

This work was supported in part by the Knight Foundation and the Office of Naval Research grant Minerva-Multi-Level Models of Covert Online Information Campaigns (N00014-21-1-2765). Additional support was provided by the Center for Computational Analysis of Social and Organizational Systems (CASOS) at Carnegie Mellon University. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Knight Foundation, Office of Naval Research, or the U.S. Government.

a life of their own [2]. These campaigns are focused on promoting pro-China rhetoric instead of content that destabilizes targeted areas. China has also used extensive bot activity to amplify messaging, such as positive coverage for the 2022 Olympics [5]. Regional analysis on the use of bots has shown that this type of inauthentic activity can perform specific roles within an information campaign [9]. Similar research on Chinese diplomatic Twitter accounts underlined China’s centralized censorship policies, but did not explore other elements of information diffusion within the network beyond the diplomatic accounts themselves [4].

The United States hosted a 2021 Summit on Democracy to address democracy-related challenges throughout the world. In response, Chinese state-sponsored Twitter accounts began a hashtag campaign around this event with anti-US sentiment. This particular event presents a time-constrained influence campaign with extensive state-sponsored support. This paper presents novel research on the following research questions to understand how agents within a CCP information campaign propagate narratives using both official state-sponsored accounts, bots, and human influencer accounts:

1. What is the nature of inauthentic and authentic accounts within this network?
2. How is information diffused through a Chinese State-Sponsored campaign?

2 Data and Methodology

2.1 Data Collection and Processing

Twitter networks consist of users and the connections that occur when a user retweets, mentions, or reply to each other. Our data collection using hashtags targets information campaigns that use Twitter’s algorithmic approach towards trending topics, which are identified and spread through users by the use of a hashtag sign preceding a topic. Using the Twitter V2 API, we collected tweets using the hashtags #WhoDefinesDemocracy and #WhatIsDemocracy, resulting in the collection of 7,798 tweets from September 1 - December 31, 2021. Our second dataset is comprised of over 200 Twitter handles of official Chinese-government affiliated Twitter accounts to include government official, government organization, and media accounts. A node in our dataset may be a Twitter agent or user, hashtag, tweet, or URL. This dataset provided 16 different meta-networks with 12,507 nodes, 79,898 edges, graph density of .0002 and a Newman Modularity score of .453, indicating a high community structure.

Twitter uses three labels for the accounts within our secondary dataset; China Government Official, China Government Organization, or China State-Affiliated Media. However, Twitter currently only labels government accounts that are heavily involved in geopolitics, state-affiliated, or are high profile in 22 countries. Due to the lack of labels across many of our state-sponsored accounts, we report the Twitter label metrics within the Results section for transparency, but make a distinction that our definition of Chinese state-sponsored accounts are any official government accounts and any Chinese media outlets.

2.2 Methods

The analytic workflow for this paper provides both user and network trends within an information campaign by a) labeling Twitter accounts for three different groups; bots, state-sponsored accounts, and all other accounts, b) conducting network analysis to understand network diffusion and echo-chamber qualities between the groups and c) dynamically analyzing the network over three periods to determine how key actors shift or persist.

Agent Labeling. We use a tier-based machine-learning tool Bothunter that classifies Twitter agents as bots using metadata and other account features [1]. To increase certainty around our bot classification for each Twitter account, we use the recommended bot probability score of .7 at which the bot classification label is most stable from flipping from one class to the other for outlying bot activity [6]. Our secondary dataset of Chinese state-sponsored accounts provides a label for agent nodes to determine which parts of our social network are state-sponsored or primarily amplify state-sponsored accounts.

Network Analysis and Echo Chambers. Using the network analysis and visualization software ORA, we use network analysis measurements such as in-degree and out-degree centralities that highlight agents disseminating versus receiving information within the network. For Twitter data, a user with high In-Degree Centrality is generally characterized with high retweet, reply, and quote frequency, whereas a user with high Out-Degree Centrality will have tweets, replies, or quotes that are frequently shared by other users. We also use community structure metrics via an E/I index; a ratio of a Twitter user’s internal and external links. We generate these values based on an agent’s behavior within a Leiden cluster group to indicate whether an actor is part of an echo-chamber in which beliefs are amplified within a group. We use the Leiden algorithm to generate these clusters based on an efficient local moving heuristic for identifying high modularity communities [10]. We use a non-parametric Kruskal-Wallis test to determine if there is a difference in average E/I index between the three actor groups based on an agent’s Leiden group clustering.

Dynamic Network Analysis. Lastly, we analyze the campaign across three periods of equal tweet density to understand how key actors and information diffusion changed throughout the duration of this campaign. This will allow us to better understand the interaction of key actors to include state-sponsored accounts, bots and other accounts across the duration of the campaign.

3 Results

Initial analysis revealed this campaign was predominantly comprised of retweets, with approximately 85% of all tweets disseminated within this network as retweets. Additionally, 54% of all tweets (original and retweets) were created by bot accounts. This section covers the Twitter language metadata analysis of the campaign and the nature of the tweets and messaging to understand bot and state-sponsored account functions for information diffusion.



Fig. 1. Top Retweets within dataset, exhibiting anti-US sentiment

3.1 Campaign Overview

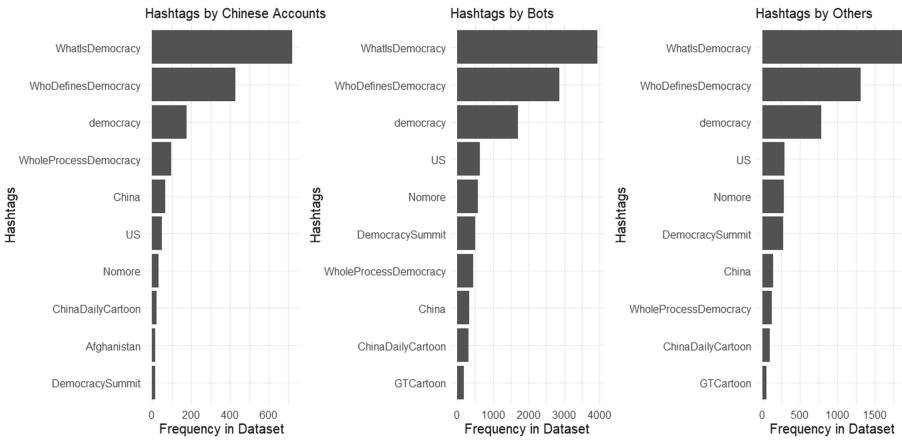
The hashtags *#whodefin democracy* and *#what is democracy* were popularized in a 2020 15-part Twitter post by a Chinese influencer on “Chinese-style Democracy” and its merits over western democracy. Prior to November 1, 2021, there were approximately 10 instances of the hashtags, all unrelated to the Democracy Summit. Chinese Diplomat to Lebanon Cao Yi posted five times in October prior to the beginning of the campaign.

The bulk of the campaign occurred within the first two weeks of December before the Democracy Summit, accounting for almost 90% of the total tweets. Using the tweet-level language metadata tags, we found 21 different languages represented in this dataset (see Table 1). The multilingual tweets indicate language diffusion targeting different populations. However, approximately 83% of our tweets are in English, indicating a clear intent to communicate to western and English-speaking audiences.

The top ten retweeted tweets accounts for nearly 25% of all tweets within our network. Of these selected tweets, three tweets discussed Chinese-style democracy and its merits in addition to promoting a Chinese-sponsored “Dialogue on Democracy” event on Chinese-style democracy. The other seven tweets contained anti-US sentiment, discussing US domestic issues such as wealth disparity, gun

Table 1. Top languages and twitter actor ratios in dataset

Language	# Tweets	% Chinese state-sponsored	% Tweets by bots
English	6520	8.2	56.6
Undetermined	324	13.5	57.1
Chinese	219	2.7	47.5
Spanish	195	18.9	51.7
Arabic	140	21.4	29.3
French	126	17.5	52.4
Russian	94	15.9	42.6

**Fig. 2.** Comparison of top Hashtags shared by each group.

violence, Black Lives Matter and racism to de-legitimize the United States (see Fig. 1). Lastly, out of the ten top tweets, only one was not created by a state-sponsored account but rather a Chinese youth outreach group.

There is extremely high correlation in hashtags used between the three different groups, indicating that bots, state-accounts and all other accounts are promoting the same messaging (see Fig. 2). By analyzing the top twenty hashtags, we found a perfect correlation in hashtags between Chinese Accounts and “Other” accounts and a .9 correlation between Chinese and Bot accounts. All three groups had the same seven hashtags that are the most widely used with the highest in-degree centrality, indicating this network contains homophilous users that promote similar content.

3.2 State-Sponsored Activity and Bot Amplification

There are 121 state-sponsored accounts within this network, accounting for approximately 3% of user accounts, but nearly half of all original tweets. These

accounts also had a 70% rate of being verified through Twitter. However, we discovered that only 17 accounts are labeled as a China Government Official or Organization, 18 accounts are labeled as state-affiliated media, and the remaining 86 accounts had no Twitter label although the account profiles indicated they were either Chinese Embassy, Ambassador, or other official government position accounts.

Of the 17 government accounts, the Chinese Ambassadors to Cuba, France, Pakistan, India, Canada, the United Kingdom, the United States, and Italy, in addition to international organizations such as the EU, UN, and ASEAN all have labels. Twitter also labels the most active and verified accounts such as China Spokepeople Lijian Zhao, Hua Chunying, and the Ministry of Foreign Affairs accounts. Twitter did not label any diplomatic accounts to African or South American countries, indicating that Twitter’s initial labeling methodology is centered on a western audience. For this reason, this paper uses the term “state-sponsored” to cover all 121 accounts within this dataset to include state-affiliated media sites and both labeled and unlabeled government accounts. The Chinese state-sponsored accounts within the network accounted for nearly 50% of original Twitter messages (see Table 2). Conversely, bot automation accounted for a smaller subset of original tweets, but the majority of retweets.

Table 2. Twitter Actor description for tweets and retweets

Actor	# Accounts	% Verified accounts	% Original tweets	% Retweets
Chinese state-sponsored	121	70	48.49	3.14
Bots activity	2038	0	16.19	61.42
Other	1420	2.1	34.6	35.44

Network Overview. A visual inspection of our network shows state-sponsored accounts at the center of Twitter user hubs (see Fig. 3). This network visualization supports the activity break-down in Table 2 pertaining to bot automation accounting for the bulk of retweet activity and indicates that the state-sponsored accounts are the main influential actors within this network whereas the bot accounts are primarily not connected to other influential nodes.

Bots and State-Sponsored accounts maintain distinct network properties within this campaign. When we separate total degree centrality into in-degree and out-degree centrality, the state-sponsored accounts generally have higher out-degree centrality paired with a lower in-degree centrality. Our dataset is filled with outliers both for bot accounts and state-sponsored accounts that create long tails within our distributions, such as the top bot account which has the highest in-degree centrality of all accounts within the dataset. Our bot retweet distribution has a long tail, with 2,268 retweets or one-third of all retweets coming from Twitter accounts that only retweeted once. This distribution stretches until we hit the outlier accounts for the top two accounts with 70 and 107 retweets.




Fig. 3. Agent x Agent Communication Network of the Twitter users, with the top “Superspreaders” labeled. Red nodes are Chinese state-sponsored accounts, green nodes are bots, and gray nodes are all other nodes. (Color figure online)

Similar to the bot accounts, a small amount of State-Sponsored accounts makes up the bulk of original tweets. Of the 711 original Chinese state-sponsored tweets, approximately 50% are from just ten accounts. Additionally, the top ten “Superspreaders” labeled in Fig. 3, or accounts with high out-degree centrality have 5,473 retweets, or account for over 70% of all tweets in this network, indicating clear centralized messaging from a small number of accounts.

Difference in Degree Distributions. If we examine the distributions for in-degree and out-degree centrality between our Bot, State-Sponsored Accounts and Other accounts in Fig. 4, the differences are clear regarding who is promoting messages and who is receiving and re-amplifying those messages. For out-degree centrality, the state-sponsored accounts have a wider interquartile range than bots and other accounts, with extreme outliers skewing within the 4th quartile. The outlying Twitter account within the “Other” category is the Chinese Youth group, which we have now seen a few times both as the creator of one of the most widely retweeted messages and a Superspreader in Fig. 3. This account behaves like a state-sponsored account due to its high out-degree centrality and amplification by both bot and other accounts. We see this same trend to a lesser degree regarding in-degree centrality, where bots have many more accounts above the interquartile range than the other two types of accounts with a large skew within the last quartile.

Difference in Echo-Chamber Qualities. We conducted a non-parametric Kruskal-Wallis test to determine if the E/I index scores for Bots, State-

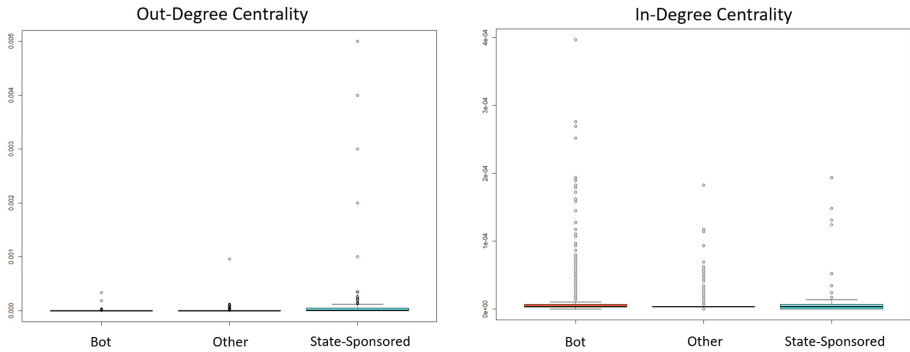


Fig. 4. In-Degree and Out-Degree distributions for State-Sponsored, Bot and Other accounts using the Agent x Agent Communication Network of the Twitter users.

Sponsored, and all other accounts are from the same distribution or if there is on average a difference in scores. We obtained a p -value $< 2.2e-16$, indicating a very statistically significant conclusion that there is a difference in average E/I index scores across the three groups. By analyzing the plot in Fig. 5, we can determine State-Actor accounts have a much lower mean value for E/I index scores, indicating that on average, these accounts participate in communities that are more prone to inner dialogue with other group members. Additionally with the mean E/I index for Bots and Other accounts within the negative value range, we can conclude this information campaign skews more towards an echo chamber environment.

Dynamic Network Analysis. We analyzed how key actors within the network shift throughout the campaign. Across the periods, top out-degree centrality accounts are dominated by state-sponsored accounts, indicating that state-sponsored accounts controlled the messaging. There was considerable consistency in the state-sponsored accounts that maintained a high volume of tweets at least one median above the average for the network across all three periods to include Cao Yi, Hua Chunying, and media account Global Times. The one account not state-sponsored is the China Youth Studio account which was active for the first two periods. Additionally, multilingual media accounts such as CGTN Arabic and CGTN Espanol were active in the second and third periods. This second period also had high out-degree centralities for the accounts to Kenya, Uganda, Cuba and the United States. Diplomatic accounts were not prominent in the third period except for the Cuban Ambassador, mainly dominated by Chinese media outlets and Spokesperson accounts.

Although top accounts for in-degree centrality are primarily bot accounts, there were two state-sponsored accounts that were in the top in-degree category one standard deviation above the mean; Diplomat Cao Yi and the spokesperson account to the United Nations. Cao Yi’s account is also the only state-sponsored account with both high in-degree and out-degree centrality measures. This account was the first account to begin tweeting prior to the campaign, indi-

cating that it played a crucial role in maintaining momentum for the duration of the campaign.




Fig. 5. Inter-quartile distribution of E/I Index scores across groups

3.3 Who Are the “Others” in This Network?

The labeling process leaves approximately 40% of accounts as “Other”, indicating they are neither a state-sponsored account nor do they have a bot probability past the .7 threshold (see Table 2). This section seeks to characterize this third group within the network that is responsible for roughly 35% of both original tweets and retweets. Due to only 125 accounts in this category contributing to original tweets, we conducted exploration on this subset to understand the nature of these accounts. We manually annotated the accounts and discovered the tweets fell into approximately three different categories; Pro-China with 132 tweets from 68 accounts, Anti-China with 30 tweets from 20 accounts, and 47 unrelated tweets from 37 accounts. The majority of the Anti-China tweets appeared to be from US-based accounts exhibiting right-wing, anti-Chinese rhetoric. The “unrelated” category contained tweets not related to China’s campaign. For example, there were many tweets regarding Myanmar and Nagaland in India that were not associated to the Summit on Democracy.

Pro-China accounts were either state-sponsored accounts that we did not have in our secondary dataset (mis-labeled “Other”), Chinese associations, influencers and regular Twitter users. The top accounts in the Pro-China group included the aforementioned Chinese youth group in addition to one more student group with 23 and 12 tweets respectively. There were 8 additional accounts that were labeled by Twitter as either a China government organization or Chinese state-affiliated media, two Russia government organization accounts and

one Cuban government account. For individual accounts that could be typically classified as “influencer” accounts, there were 20 accounts by Chinese individuals that typically retweeted many state-sponsored tweets, in addition to about 14 accounts by Twitter users in Ethiopia. This may indicate that apart from potential Chinese diaspora Twitter users, there is also a small contingent of users from other countries (particularly in Africa) participating.

4 Conclusion

This work analyzed a limited information campaign by the Chinese government around the 2021 Democracy Summit, concluding key differences in network measures around State-Sponsored accounts, bots and all other accounts in addition to identifying influencer accounts that behave like state-sponsored accounts. This campaign exhibited strong centralized control of the narrative propagation by select CCP spokespeople and news agencies to promote a distinctly anti-US campaign, with message dissemination to different target populations. A limitation to consider is that our state-sponsored accounts may not capture the entire ground truth as our supplementary data on state-sponsored Twitter accounts is manually created and likely under-capturing the extent of official state-sponsored accounts within the network.

References

1. Beskow, D.M., Carley, K.M.: Bot-hunter: a tiered approach to detecting & characterizing automated activity on twitter. In: Conference Paper. SBP-BRIMS: International Conference on Social Computing, Behavioral-Cultural Modeling and Prediction and Behavior Representation in Modeling and Simulation, vol. 3, p. 3 (2018)
2. Beskow, D.M., Carley, K.M.: Characterization and comparison of Russian and Chinese disinformation campaigns. In: Shu, K., Wang, S., Lee, D., Liu, H. (eds.) *Disinformation, Misinformation, and Fake News in Social Media*. LNCS, pp. 63–81. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-42699-6_4
3. DiResta, R., Goldstein, J.A.: China’s fake twitter accounts are tweeting into the void. *Foreign Policy* (2021)
4. Huang, Z.A., Wang, R.: Building a network to “tell china stories well”: Chinese diplomatic communication strategies on twitter. *Int. J. Commun.* **13**, 24 (2019)
5. Myers, S.L., Mozur, P., Kao, J.: How Bots and Fake Accounts Push China’s Vision of Winter Olympic Wonderland. *ProPublica* (2022)
6. Ng, L.H.X., Robertson, D.C., Carley, K.M.: Stabilizing a supervised bot detection algorithm: how much data is needed for consistent predictions? *Online Soc. Networks Media* **28**, 100198 (2022)
7. Schafer, B.: China Fires Back at Biden with Conspiracy Theories About Maryland Lab. *Foreign Policy* (2021)
8. TwitterSafety: Disclosing state-linked information operations we’ve removed (2021). https://blog.twitter.com/en_us/topics/company/2021/disclosing-state-linked-information-operations-we-ve-removed

9. Uyheng, J., Carley, K.M.: Characterizing bot networks on twitter: an empirical analysis of contentious issues in the Asia-Pacific. In: Thomson, R., Bisgin, H., Dancy, C., Hyder, A. (eds.) SBP-BRiMS 2019. LNCS, vol. 11549, pp. 153–162. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-21741-9_16
10. Waltman, L., van Eck, N.J.: A smart local moving algorithm for large-scale modularity-based community detection. *Eur. Phys. J. B* **86**(11), 1–14 (2013). <https://doi.org/10.1140/epjb/e2013-40829-0>