Computational and Mathematical Organization Theory
https://doi.org/10.1007/s10588-023-09380-9

ORIGINAL PAPER

#WhatIsDemocracy: finding key actors in a Chinese influence campaign

Charity S. Jacobs[1] · **Kathleen M. Carley**[1]

Accepted: 15 June 2023

Abstract
The rapid increase in China's outward digital presence on western social media platforms highlights China's priorities for promoting pro-Chinese narratives and stories in recent years. Simultaneously, China has increasingly been accused of launching information operations using bot activity, puppet accounts, and other inauthentic activity to amplify its 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. China uses layers of state-sponsored accounts, bots, and non-bot accounts to promote its messaging. Lastly, we describe how China uses localized campaigns under a more extensive umbrella campaign for information diffusion toward targeted audiences.

Keywords China · Influence operations · Network analysis · Bots · Social media

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 using media with an international audience, passing domestic legislature as recently as 2007 to increase the country's discourse power by creating a Chinese-controlled media platform Xinhua. Xinhua and other similar state media platforms, such as the Chinese Global Television Network (CGTN), can shape narratives about China at the global level (DiResta and Goldstein 2021). While China's iron-clad censorship policies, such

✉ Charity S. Jacobs
csjacobs@cmu.edu

Kathleen M. Carley
kathleen.carley@cs.cmu.edu

[1] Software and Societal Systems, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh 15213, PA, USA

Published online: 01 July 2023

as the Great Firewall to filter and block content internally, are well-known (Crandall et al. 2007; King et al. 2013), how China shapes international dialogue is less known. Increasingly, Chinese Communist Party (CCP) officials are using Western social media platforms to defend Chinese national interests as part of its continued discourse efforts to shape international opinion. The use of these official accounts is not by themselves unique from any other country's use of social media. However, these accounts are increasingly scrutinized for spreading disinformation and conspiracy theories such as the origins of COVID-19, human rights abuse in Xinjiang, and bot amplification around the 2022 Beijing Olympics (Feng 2022; Schafer 2021; Twitter Safety 2021). As China continues using social media platforms for messaging international audiences, understanding its overt social media activity may help researchers identify inauthentic activity patterns that are not easily attributable.

The United States hosted a 2-day "Summit on Democracy" event in mid-December 2021 to address democracy-related challenges worldwide, inviting over one hundred countries to participate. In response, Chinese state-sponsored Twitter accounts began a hashtag campaign around this event with anti-US sentiment in addition to messaging for its internal event called "Dialogue on Democracy" on December 2, 2021. China's hashtag campaign around the Summit on Democracy represents a time-constrained influence campaign with extensive state-sponsored support. This paper presents network methods to identify and understand how key actors within a Chinese information campaign propagate narratives on a social media platform. Specifically, we explore the following:

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 Related work

State-sponsored activity on social media platforms such as Twitter has gained tremendous momentum recently, with evidence of organized campaigns found in over 80 countries (Bradshaw et al. 2021). These campaigns imitate human behavior to persuade, manipulate, coerce, or crowd out targeted online audiences through information tactics such as distorting online discussion, bridging together, and exciting users towards a viewpoint or topic (Beskow and Carley 2019).

Information operations may use diverse methods to manipulate the information domain, including the use of bots controlled by automation, troll farms that use professionalized groups to conduct a coordinated activity, cyborg accounts that may switch between software and human control, and puppet accounts where a person may use a fictitious persona to blend in with a community. Propaganda and disinformation efforts that use these methods to imitate human social media interactions and manipulate discussion have been shown to be effective at shaping public opinion (Bradshaw et al. 2021; Woolley and Howard 2017). Bot accounts in particular, are effective at influencing public opinion and narratives due to the cost efficiency, paired with social media users' poor ability to determine if another account is

authentic or a bot (Everett et al. 2016). In short, the modern state actor with few ethical qualms maintains a swiss army knife of options to shape the discourse of ideas and public opinion through information maneuvers and tactics waged in the information domain. These tactics might include using any of the inauthentic account types, as mentioned earlier, to shape both narratives and the online communities where those narratives are propagated (Beskow and Carley 2019; Blane et al. 2022).

Inauthentic activity is often comprised of the following; bots which are accounts controlled by automation, troll farms that use professionalized groups coordinating activity, and puppet accounts that use fictitious personas. These, among others, have been well-documented methods for disseminating information and disinformation by imitating human social media interactions (McCombie et al. 2020; Hwang et al. 2012; Ferrara et al. 2016; Shu et al. 2017). Bots effectively change public opinion and shift narratives on social media, creating a clear incentive for bad actors to manipulate the information domain (Bessi and Ferrara 2016). Additionally, bots are useful as an amplification device in spreading and disseminating information campaigns throughout a social network more efficiently than humans (Woolley and Guilbeault 2017). Within the Twitter domain, tweets are easily manipulated via the retweeting function, a method easily exploited by bot automation (Boyd et al. 2010). While Twitter periodically takes down bot networks, especially those associated with state-sponsored activity, bots constantly adapt to circumvent bot detection measures (Luceri et al. 2019).

Much of what is known about Chinese social media campaigns is derived from social platforms releasing data on inauthentic activity network takedowns, observational studies such as this paper, and qualitative research based on Chinese original government documents. China's information campaigns have typically promoted pro-China rhetoric and attacked perceived state enemies (Uren et al. 2019; Phillips et al. 2022). Additionally, China has traditionally avoided using memes that often take on a life of their own and potentially undermine centralized messaging control (Beskow and Carley 2020). Extensive computational journalism by media outlets like the New York Times, ProPublica, and the Australian Strategic Policy Institute has coordinated with Twitter in analyzing these networks. Key findings include evidence of the Chinese government outsourcing Twitter campaigns to private companies, the use of a layered ecosystem of state-sponsored accounts and influencers/foreign voices, and the prolific use of fake accounts for amplification (Ryan et al. 2021; Mozur et al. 2021; Myers et al. 2022).

Past research, specifically on Chinese diplomatic Twitter accounts, determined that information diffusion occurred primarily from key state-sponsored outlets as a function of China's centralized censorship policies. However, this research did not explore other elements of information diffusion within the network beyond state-sponsored Twitter accounts (Huang and Wang 2019). Additionally, this research was conducted when there were only 14 state-sponsored accounts compared to the present count of over 300 accounts.

This paper contributes to current research by applying a network mapping to China's state-sponsored social media network to identify and understand how China amplifies its discourse. A network approach allows us to empirically quantify the different layers of this messaging ecosystem in terms of individuals and types of

actors, including bots and influencers, in addition to understanding the interactions between these actors. Additionally, we use network metrics to identify other country-state actors and further investigate smaller, more localized information campaigns that fall under an international campaign.

3 Methods

3.1 Data collection

Twitter networks consist of users and the connections that occur when users retweet, mention, 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 with 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 7798 tweets from September 1–December 31, 2021. Our second dataset contains over 300 Twitter handles of official Chinese-government affiliated Twitter accounts, including government official, government organization, and media accounts.

3.2 Terminology and ethics

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 heavily involved in geopolitics, state-affiliated, or high profile in 22 countries. Due to the lack of labels across approximately 70% of our state-sponsored accounts, we report the Twitter label metrics within Sect. 4 for transparency. We make a distinction that our definition of Chinese state-sponsored accounts are any official government accounts and any Chinese media outlets.

We label and describe state-sponsored accounts by an individual's name, title, or media name. This labeling process includes accounts representing the Chinese government and other state-sponsored participants within this dataset. All other accounts are anonymized when described for privacy purposes.

3.3 Mixed-methods pipeline

The analytic workflow (see Fig. 1) 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, agent properties and echo-chamber qualities between the groups and (c) dynamically analyzing the network over three periods to determine how key actors shift or persist. We use a mixed-methods pipeline using a sequential explanatory strategy to first empirically define

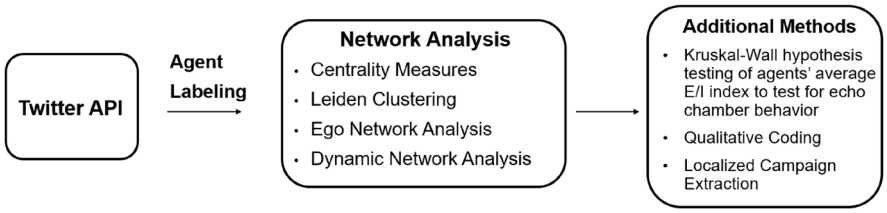

Fig. 1 Mixed methods pipeline for network analysis and campaign exploration

our network data and its agents, with follow-on methods to qualitatively analyze the messaging within this campaign.

Agent Labeling We use a tier-based machine-learning tool, Bothunter, that classifies Twitter agents as bots using metadata and other account features (Beskow and Carley 2018). To increase certainty around our bot classification for each Twitter account, we use the recommended bot probability score of 0.7, at which the bot classification label is most stable from flipping from one class to the other for outlying bot activity (Ng et al. 2022). 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 characterize agents disseminating and receiving information within the network. For Twitter data, a user with high In-Degree Centrality is generally characterized by a high retweet, reply, and quote frequency. In contrast, users with high Out-Degree Centrality will have tweets, replies, or quotes that other users frequently share.

We also use community structure metrics via an external-internal or E/I index (Krackhardt and Stern 1988), a Twitter user's internal and external links ratio. 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 (Waltman 2013). We use a non-parametric Kruskal-Wallis test to determine if there is a difference in the average E/I index between the three actor groups based on an agent's Leiden group clustering.

We generate ego networks that encompass an individual or ego along with its peers or alters and provide unique structural properties for exploring information diffusion among our outlier accounts (Arnaboldi et al. 2017). We use each group's top ten tweet creators (State, Bot, and Other) to generate three ego networks to describe better-shared properties and agent connections. Due to their high activity, the selected accounts represent the top outlier accounts for each category. By examining the surrounding networks of agents for the outlying state-sponsored accounts, bots, and other accounts, these networks better inform us of the role functions within the campaign.

Qualitative Coding We use an iterative process to qualitatively group themes together by analyzing a data sequence and determining high-level patterns for the appropriate categories. This method is grounded in inductive or ground-up coding where categories are derived based on observations from the data. For all areas of analysis where coding is required, we reviewed approximately 20% of the data to determine high-level themes and coded the data using the derived codes. We use an inter-coder agreement via Cohen's kappa score to determine the quality of our independently coded data.

Finding localized campaigns Our data in its entirety represents an international campaign. We propose extracting sub-networks that may represent localized campaigns conducted through China's Consulate offices based on research regarding China's past information operations (Kania 2019; Jacobs and Carley 2022). By using our known labeled agents from the labeling process, we can find similar unlabeled agents by looking at outliers within our closeness centrality network that determines the closeness of a node to other nodes within a network. We examine the nodes at least one standard deviation above the mean and explore these accounts manually to determine if they represent other non-Chinese state-actor accounts. We also use the tweet language metadata tags to add tweets if the targeted audience uses a specific language. We extract the sub meta-network of all known state actors and follow-on tweets labeled for a given language.

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 this campaign. Dynamic Network Analysis will allow us to better understand the interaction of key actors to include state-sponsored accounts, bots and other accounts across the campaign's duration.

4 Results

Initial analysis revealed this campaign was predominantly comprised of retweets, with approximately 85% of all tweets (about 6686 tweets) disseminated as retweets that forward an original tweet. This section covers the Twitter language metadata analysis of the campaign, qualitative categorization of the State-Sponsored tweets, network characterization, and localized information diffusion used by China's social media ecosystem to amplify state-sponsored information campaigns.

4.1 Campaign overview

The hashtags #whodefinesdemocracy and #whatisdemocracy 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.

Table 1 Top languages and Twitter actor ratios in dataset

Language	# Tweets	% Chinese state accounts	% 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

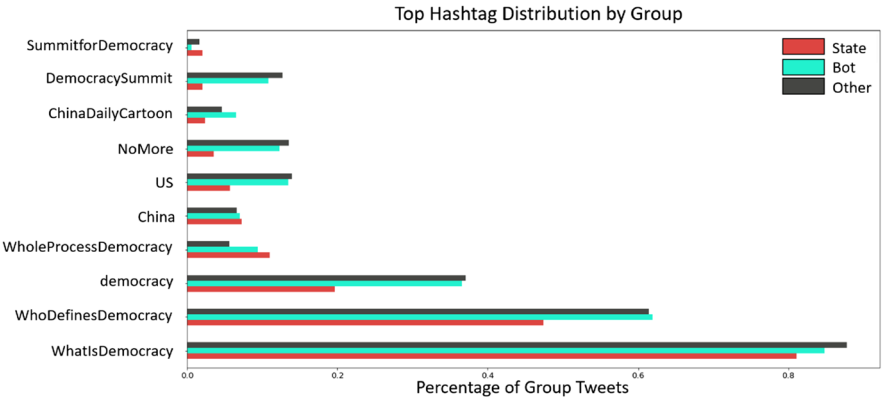

Fig. 2 Comparison of top Hashtags shared by each group. Bots and Other accounts have higher rates of using top hashtags due to the retweet nature of the network. State accounts create more original tweets, which have slightly less use of top hashtags

The bulk of the campaign occurred within the first 2 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.

We compared hashtag distributions between types of actors to determine if bots use certain hashtags more than other accounts. There is an 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). The hashtag distributions indicate that Bots and Other accounts use top hashtags at a slightly higher rate than state accounts. This is a byproduct of the retweeting function for this network, where retweets are more likely to have hashtags and original tweets with no hashtags are less likely to be retweeted. By analyzing the top twenty hashtags, we found a perfect correlation in hashtags between Chinese Accounts and "Other" accounts and a 0.9 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.

To better categorize this campaign's themes, we analyzed the top 100 tweets. When included with subsequent retweets, these tweets accounted for nearly 4865 tweets or approximately 62% of all tweets. By examining the first 20 tweets, we determined three categories of tweets: Pro-China, Anti-US, and general tweets on Democracy. We then re-coded all 100 tweets independently, resulting in an initial Cohen's kappa of 0.87 for a raw agreement on 94% of tweets, indicating a strong agreement between the graders. We then discussed and jointly labeled the tweets without agreement (Fig. 3).

China released an official statement through its Ministry of Foreign Affairs page on the US Summit for Democracy and highlighted three main points listed verbatim below (Spokesperson's Remarks 2021).

1. The US is not a "beacon of democracy", and the American-style democracy has deviated from the essence of democracy.
2. A country's path to democracy should be chosen independently by its own people, rather than imposed from outside.
3. Stoking division and confrontation in the name of democracy is to backpedal in history, and will bring nothing but turmoil and disaster to the world.

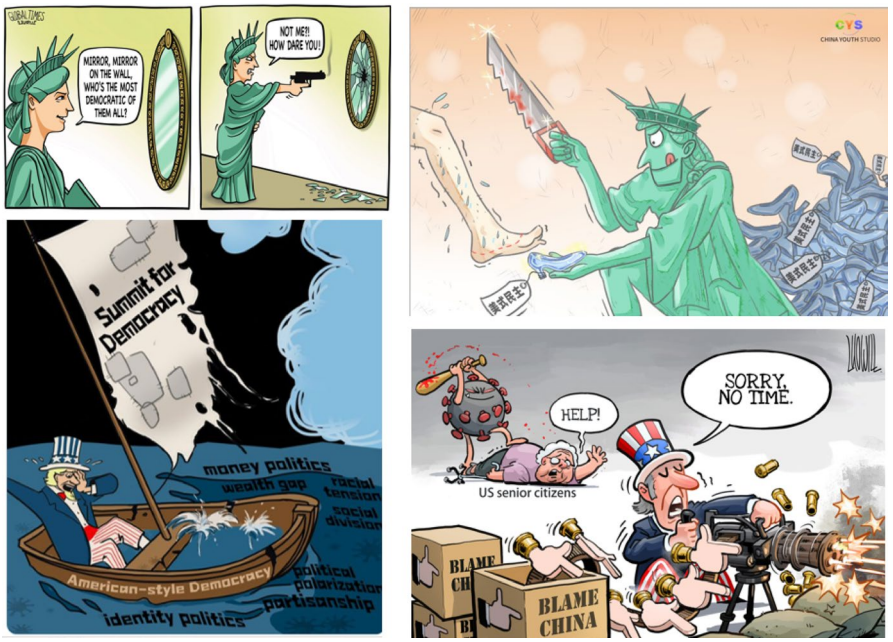

Fig. 3 Example Cartoons in anti-US tweets. The Anti-US tweet category contained many cartoons and video links, utilizing multi-modal sources of content. These cartoons were extracted from tweets by the Ambassador to Lebanon (Yi 2021), "Youth Group", (China Daily 2021), and (Global Times 2021)

These official statements validate our three topic categories. Except for the United States, China does not name any other countries in this campaign, with the exception of hashtags associated with past US conflict zones such as Afghanistan, Iraq, and Syria. China is primarily positioning itself as a foil and prime opposing power to the United States, even within a constrained information campaign. Tweets within this campaign use well-controlled messaging to support critical points the Chinese government seeks to underline.

4.2 State-sponsored activity and bot amplification

An overview of the types of actors and their contributions to the campaign indicates that this campaign was driven primarily by state accounts (see Table 2). There are 137 state-sponsored accounts within this network, accounting for only 3% of user accounts but over half of all original tweets. Conversely, bot automation accounted for a smaller subset of original tweets but the majority of retweets. The state accounts also had a high rate of being verified through Twitter. However, we discovered that only 17 accounts were labeled as a China Government Official or Organization, 18 accounts were labeled as state-affiliated media, and the remaining 102 accounts had no Twitter label. These unlabeled accounts contained profiles indicating 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, and international organizations such as the EU, UN, and ASEAN all have labels. Twitter labels the most active and verified accounts, such as China Spokespeople Lijian Zhao, Hua Chunying, and the Ministry of Foreign Affairs accounts. Twitter did not label any diplomatic accounts in African or South American countries, indicating that Twitter's initial labeling methodology focuses on a western audience. For this reason, this paper uses the term "state-sponsored" to cover all state accounts within this dataset, including state-affiliated media sites and both labeled and unlabeled government accounts.

Network overview A visual inspection of our network shows state-sponsored accounts at the center of Twitter user hubs (see Fig. 4). 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

Table 2 Twitter Actor description for tweets and retweets

Actor	# Accounts	% Verified accounts	% Original tweets	% Retweets
Chinese state-sponsored	137	64.0	64.7	3.9
Bots activity	2037	0	13.6	67.0
Other	1380	1.3	34.6	34.6

China's state-sponsored accounts account for the majority of original tweets, while Bots are the primary amplification actor for original tweets

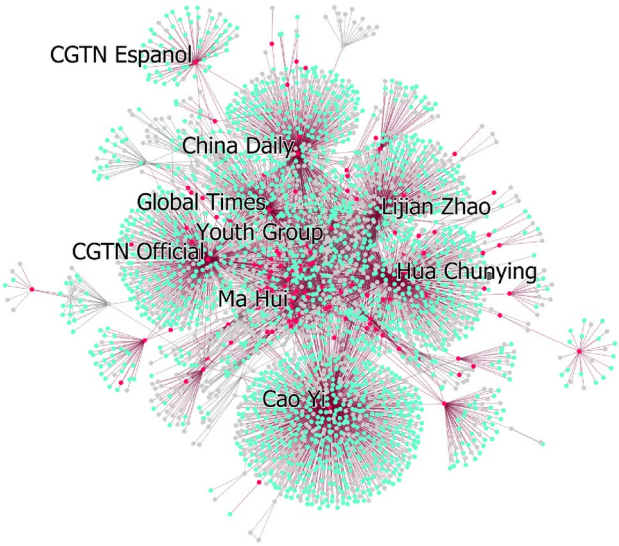

Fig. 4 Agent × Agent Communication Network of the Twitter users, with the top "Superspreaders" labeled. Red nodes are Chinese state-sponsored accounts, green nodes are bots, and gray nodes are all other nodes. All Superspreaders were state-sponsored accounts except for our "Youth Group" account. (Color figure online)

the main influential actors within this network. The bot accounts are primarily used to retweet original tweets by state actors as amplifiers.

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 shaped by 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 2268 retweets or one-third of all retweets 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. Similar to the bot accounts, a small amount of State-Sponsored accounts make 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. 4, or accounts with high out-degree centrality, have 5473 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. 5, the differences are apparent 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

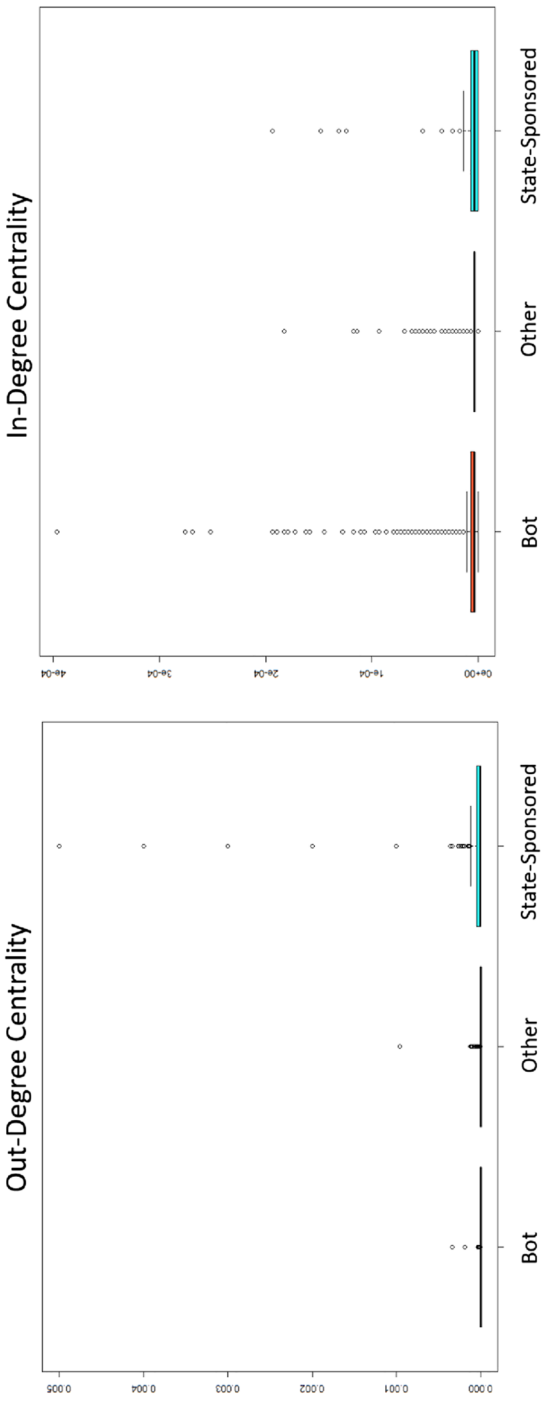

Fig. 5 In-degree and out-degree distributions for state-sponsored, bot and other accounts using the Agent × Agent Communication Network of the Twitter users

other accounts, with extreme outliers skewing within the fourth quartile. The outlying Twitter account within the "Other" category is a Chinese Youth organization that we will call "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. 4. 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 significant 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-Sponsored accounts, 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 statistically significant conclusion that the average E/I index scores differ across the three groups. By analyzing the plot in Fig. 6, we can determine that State-Actor accounts have a lower mean value for E/I index scores, indicating that, on average, these accounts participate in communities that are prone to inner dialogue with other group members. Additionally, we can conclude that this information campaign skews towards an echo chamber environment because all groups result in a negative mean E/I value.

Ego network properties Due to the high number of outliers within our network, we examine the ego networks for the three different actor types to determine if one group has more ties or shares more alter nodes. We conducted an ego network comparison by creating three ego networks for each actor type using the top ten accounts for each category with the highest number of tweets, representing outliers for both our in-degree centrality and out-degree centrality metrics. Our results in Table 3 highlight how the top state-actor accounts had more significant variance in mean connected alter nodes and links, whereas bots had the lowest variance.

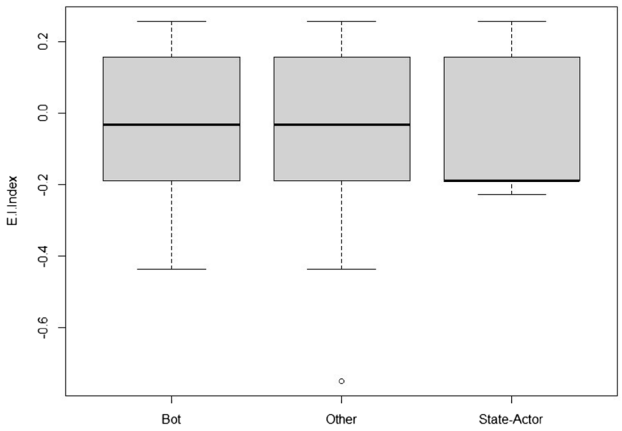

Fig. 6 Inter-quartile distribution of E/I index scores across groups

Table 3 Ego network comparison between actor types

Actor	Ego network statistics[a]			Mean ego network[b]		
	Nodes	Links	Density	Nodes	Links	Density
State Accounts	594 (713.8)	5106 (7065)	0.02 (0.02)	14	61	0.19
Bots	155.5 (43.8)	857 (217.5)	0.04 (0.01)	45	303	0.23
Other	107.4 (134.3)	582.4 (815.4)	0.09 (0.04)	18	65	0.50

Within our mean ego networks, Bots shared the most Chinese state-sponsored alter nodes, indicating the top bot accounts are amplifying the same state-sponsored accounts

[a]Mean (standard deviation)

[b]Mean ego network contains the nodes and links that occur in at least half of the individual ego networks

By examining the mean ego network where only nodes and links in at least 50% of a group's ego networks are examined, state accounts had the lowest shared nodes, links, and density. Bots shared the most nodes and links, indicating more shared alter nodes. When we examine shared agent nodes within our mean ego networks, we find that our bot ego network has nine State Accounts and one other account "Youth Group". This finding indicates that top bots all primarily amplify the same state-sponsored accounts. For state-sponsored accounts, we see three bot accounts and three state accounts. These Chinese accounts are also part of our Superspreader group and indicate that even in the top ten State Accounts, the lower-activity accounts amplify the higher-activity accounts. Lastly, our "Other" mean ego network only contained two state-actor accounts, potentially indicating less coordinated activity between the top actors in this group and the campaign itself.

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 throughout the campaign. 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, including Cao Yi, Hua Chunying, and the media account Global Times. The one account that is not state-sponsored is "Youth Group", which was active for the first two periods. Additionally, multilingual media accounts such as CGTN Arabic and CGTN Español were active in the second and third periods. This second period also had high out-degree centralities for the consulate 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, two state-sponsored accounts 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 high in-degree and out-degree centrality measures. This account was the first to

begin tweeting before the campaign, indicating that it played a crucial role in maintaining momentum for the campaign's duration.

4.3 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 have a bot probability past the 0.7 threshold (see Table 2). This section seeks to characterize this third group within the network responsible for roughly 35% of original tweets and retweets. Due to only 125 accounts in this category contributing to original tweets, we explored 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. Most 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, many tweets regarding Myanmar and Nagaland in India were not associated with the Summit on Democracy.

Pro-China accounts were state-sponsored accounts that we did not have in our secondary dataset (mislabeled "Other"), Chinese associations, influencers, and regular Twitter users. The top accounts in the Pro-China group included the aforementioned "Youth Group" and one more student group "Youth Group 2" with 23 and 12 tweets, respectively. There were eight additional accounts that Twitter labeled either a China government organization or Chinese state-affiliated media, two Russian 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 finding may indicate that apart from potential Chinese diaspora Twitter users, a small contingent of users from other countries (particularly in Africa) are also participating.

4.4 Localized information operation campaigns

Using the closeness centrality measurement that determines the closeness of a node to other nodes within a network, we examined the nodes that are at least one standard deviation above the mean. This section discloses the non-Chinese state-sponsored

Table 4 Other countries and local Chinese activity

Country	No. accounts	No. tweets	No. tweets by Chinese accounts	No. retweets for Chinese tweets
Cuba	11	11	28	129
Uganda	7	8	17	44
Russia	4	2	38	56

China's largest engagement was with Cuban accounts

activity we found in Table 4 along with the interactions between these countries and their Chinese Diplomatic Representatives. Additionally, we examine language-specific campaigns with no geographic boundary or government representation.

Uganda: promoting a local symposium on democracy Uganda represents an interesting example of China taking strategic advantage of another country's weakened ties with the United States to strengthen its ties. It is important to note that the United States did not invite Uganda to the Summit on Democracy. This fact and the adverse reaction from Uganda's Ministry of Foreign Affairs were widely reported in Uganda news outlets (Mufumba 2021). We found tweets advertising a December 10th Democracy Symposium organized by Uganda's Kampala International University and the Uganda Council on Foreign Relations under the theme "Dissecting Western Democracy in Africa," with China providing a keynote speech.

Leading up to Uganda's local symposium on democracy, we see limited amplification of a Chinese state tweet on December 8, which advertised an article by the Chinese Ambassador to Uganda, stressing the similarities between Uganda and China and the potential for partnerships. This article did not mention the US by name but mentioned that a "certain country is putting together the so-called Democracy Summit as self-styled leader of democracy (Zhang 2021)." Of note, the Russian Ambassador to Uganda also retweeted this tweet. We see coordination via retweets and mentions between official Uganda accounts, the Ugandan Kampala University, and Chinese Diplomats to Uganda to promote a joint symposium on democracy while also using anti-US rhetoric (see Fig. 7).

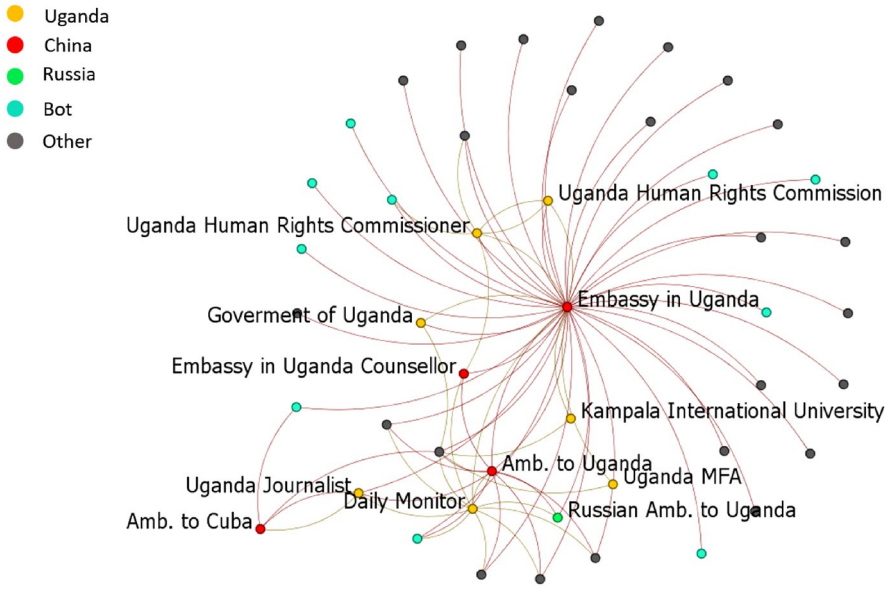

Fig. 7 Agent × Agent Communication Network of the Twitter users in Uganda's localized campaign. This network represents the coordinating partners to amplify messaging against the US Summit on Democracy and boost Uganda's Symposium on Democracy

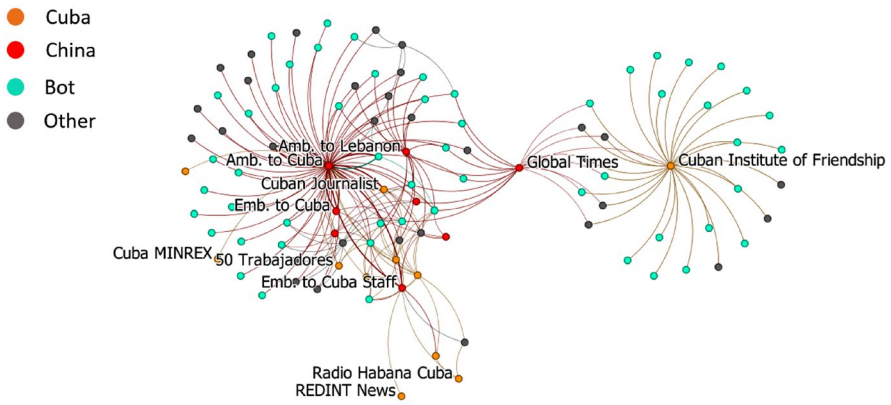

Fig. 8 Agent × Agent Communication Network of the Twitter users in Cuba's localized campaign. China's Ambassador to Cuba and the Cuban Institute of Freedom were at the centers of gravity, bridged by China's Global Times outlet

Cuba: integrating with local Cuban voices Chinese Diplomats to Cuba coordinated with Cuban Officials in addition to the Cuban Institute of Friendship to prop up a Cuba-facing localized campaign (see Fig. 8). We found tweets advertising an article the Chinese Ambassador to Cuba wrote within a local Cuban newspaper on China's relationship to Cuba on the subject of democracy. This article mirrored much of the verbiage as the Ambassador to Uganda's article such as the line "... but certain countries call themselves 'democratic leaders' and convene the so-called 'Summit for Democracy' to categorize the countries of the world and label them as 'democratic' and 'undemocratic'."

China's Cuban messaging centers around China's Ambassador to Cuba Ma Hui. His tweets and subsequent amplification account for approximately 63% of tweets. Within this subnetwork, China's consulate staff use a network of Cuban journalists and local news outlets to promote its messaging. Most other tweets in this network originate from and are amplified by the Cuban Institute of Friendship for the People Asia Pacific Department. China's Global Times news network operates as a bridging agent, providing messaging that both the Ambassador and the Cuban Institute retweet and amplify to the Cuban audience.

Russia: limited activity There was limited participation in this dataset from Russia, with only four Russian state-sponsored accounts. There was cooperation to amplify an RT (formerly Russia Today) story on Julian Assange between a Russian Journalist and Chinese state accounts retweeting the story. These tweets used the hashtags #FreeAssangeNOW and #WhatisDemocracy within Russian-language tweets. However, we found no evidence of widespread coordination between Russian and Chinese accounts about a specific localized event. We primarily found all users in this network, including the Russian accounts retweeting CGTN Russia tweets. Russian state accounts appear to be consumers and propagators of Chinese state-affiliated media outlets designed for the Russian audience.

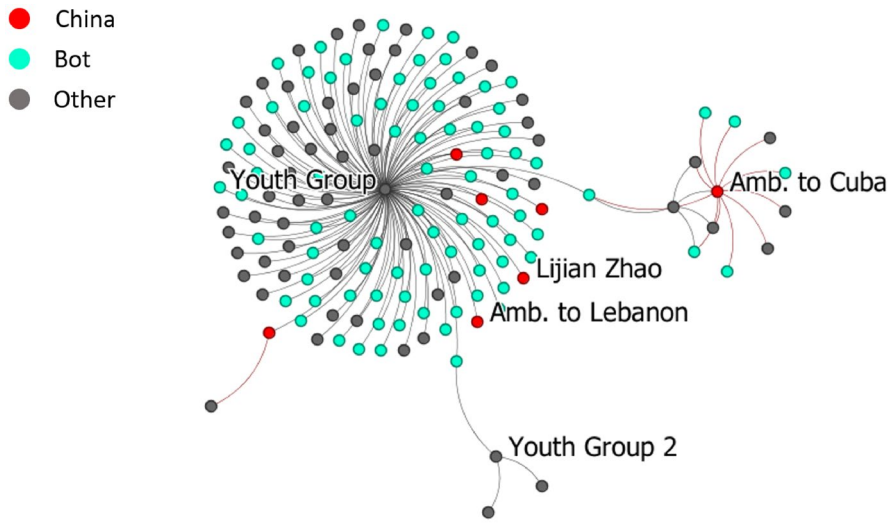

Fig. 9 Agent × Agent Communication Network of Twitter Users associated with Chinese language tweets. For this campaign, China primarily used one non-state-sponsored account "Youth Group" to create original tweets that are then amplified by bots, state accounts and other accounts

Chinese language tweets dominated by non-state-account We examined whether there was different messaging specific to Chinese language speakers, potentially targeting the mainland Chinese diaspora. Our analysis of Chinese language tweets reveals that approximately 80% are amplifying one specific non-state-sponsored account that we refer to as "Youth Group" (see Fig. 9). Our prior sections revealed that this account appears like a state account from a network perspective due to the amplification of original tweets by both bots and state accounts. This account also uses cartoons that appear to be original content, not taken from other sites like CGTN or Global Times. Another youth group which we call "Youth Group 2" also creates original content with the use of meme-like graphics and state-sponsored video-embedded links, but does not have the same scale of amplification as Youth Group. Lastly, the Ambassador to Cuba Ma Hui is also in this network, with Chinese language tweets amplifying the Cuban Ministry of Foreign Affairs negative statements towards the Summit on Democracy. However, his tweets are not amplified by other state accounts and so are not likely a key component for messaging to a Chinese audience.

Overall themes Overall, we see a few patterns develop in localized campaigns. Chinese Ambassadors mirrored official messaging for localized messages to countries. Many Ambassadors used local news outlets, such as the Ambassadors to Cuba, Uganda, and Saudi Arabia, to better reach local audiences. Twitter is then used to propagate lengthier Op-Ed pieces by China's Ambassadors. Additionally, including local influencers through interviews, symposiums, and discussions highlights China's inclusiveness towards these countries while still

promoting Anti-US rhetoric. Chinese-language tweets are an outlier in localized messaging, where one non-state-sponsored account generates most tweets that are amplified. State accounts also amplify this account, indicating potentially some control or endorsement over the messaging.

5 Discussion

In this work, we analyzed a limited information campaign by the Chinese government around the 2021 Democracy Summit. Key differences in network measurements around State-Sponsored accounts, bots and all other accounts in allow researchers to potentially identify Chinese information campaigns from the agents and how they interact alone. Our state-sponsored accounts were the predominant superspreaders for this campaign, maintaining high outlying out-degree centrality while the most active bots amplified this messaging with high in-degree centrality, indicating high retweet activity. By analyzing superspreaders in this campaign, we also identified non-state-sponsored accounts that behave like state-sponsored accounts, specifically in the creation and amplification of original content. This campaign exhibited strong centralized control of the narrative by select CCP spokespeople, Chinese news agencies, and influencer accounts such as Youth Group to promote a distinctly anti-US campaign. These efforts were aided by amplification from a layered ecosystem of other less prominent Chinese state accounts, bots, and other accounts.

Additionally, we examined China's localized campaign strategies to target different populations, primarily through China's diplomatic offices. These campaigns were frequently in coordination with local government, news outlets, and community members. While the impact of these campaigns was not measured, we found evidence of China using similar tactics in different countries to disseminate Chinese consulate messages with similar verbiage and rhetoric to local populations. This indicates a potentially global strategy with localized sub-components for diffusing messaging both at the international and local levels. China was able to better exploit and promote its anti-US rhetoric by using a country's grievances towards the US, all of whom were excluded from the US Summit on Democracy.

6 Conclusions and future work

Our work contributes to the growing body of research on state-sponsored activity within the social media realm by using network metrics to characterize different roles within a Chinese information campaign. Additionally, understanding how state actors spread information campaigns internationally and whom they target with localized messaging enables more transparency for social media consumers. Due to the growing use of a state-sponsored activity to manipulate public opinion, this research will grow in importance as social media platforms and researchers struggle to detect emerging inauthentic behavior that may adapt over time.

A limitation to consider is that our state 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 extenst of official state-sponsored accounts within the network. Future work would benefit from programmatically labeling all state and media accounts based on profile features, thereby increasing transparency around state-sponsored behavior worldwide.

Acknowledgements We would like to thank Lynnette Ng and Joshua Uyheng for their invaluable inputs to this paper. 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.

Funding Open Access funding provided by Carnegie Mellon University.

References

Arnaboldi V, Conti M, Passarella A, Dunbar R (2017) Online social networks and information diffusion: the role of ego networks. Online Soc Netw Media 1:44–55. https://doi.org/10.1016/j.osnem.2017.04.001

Beskow DM, Carley KM (2018) Introducing bothunter: a tiered approach to detection and characterizing automated activity on twitter. In: Bisgin H, Hyder A, Dancy C, Thomson R (eds) International conference on social computing. Behavioral-cultural modeling and prediction and behavior representation in modeling and simulation. Springer

Beskow DM, Carley KM (2019) Social cybersecurity: an emerging national security requirement. Mil. Rev. 99(2):117

Beskow DM, Carley KM (2020) 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. Lecture notes in social networks. Springer, Cham. https://doi.org/10.1007/978-3-030-42699-6_4

Bessi A, Ferrara E (2016) Social bots distort the 2016 US Presidential Election Online Discussion, vol 21, Number 11. First Monday. https://ssrn.com/abstract=2982233

Blane J, Bellutta D, Carley KM (2022) Social-cyber maneuvers during the COVID-19 vaccine initial rollout: content analysis of tweets. J Med Internet Res 24(3):e34040

Boyd D, Golder S, Lotan G (2010) Tweet, tweet, retweet: conversational aspects of retweeting on twitter. In 2010 43rd Hawaii international conference on system sciences. IEEE, (pp. 1–10)

Bradshaw S, Bailey H, Howard P (2021) Industrialized disinformation: 2020 global inventory of organized social media manipulation. Computational Propaganda Research Project. https://demtech.oii.ox.ac.uk/research/posts/industrialized-disinformation/

China Daily (2021) Double standards on democracy and human rights #WhatIsDemocracy #ChinaDailyCartoon #US [Tweet]. Retrieved from https://twitter.com/ChinaDaily/status/1472159711357456386

Crandall J, Zinn D, Byrd M, Barr E, East R (2007) ConceptDoppler: a weather tracker for Internet censorship. CCS 7:352–365

DiResta R, Goldstein J (2021) China's Fake Twitter accounts are tweeting into the void. foreign policy. https://foreignpolicy.com/2021/12/15/china-twitter-trolls-ccp-influence-operations-astroturfing/. Accessed 1 July 2022

Everett R, Nurse J, Erola A (2016) The anatomy of online deception: what makes automated text convincing? In: Proceedings of the 31st annual ACM symposium on applied computing, pp 1115–1120

Feng J (2022) China spreading Xinjiang disinformation on 'Global Scale': State Department. Newsweek. https://www.newsweek.com/china-disinformation-xinjiang-uyghur-human-rights-state-department-1736471 Accessed 1 July 2022

Ferrara E, Varol O, Davis C, Menczer F, Flammini A (2016) The rise of social bots. Commun ACM 59(7):96–104. https://doi.org/10.1145/2818717

Global Times (2021) #GTCartoon: the statue of liberty loses it after a mirror tells the truth. #WhoDefinesDemocracy [Tweet]. Retrieved from https://twitter.com/globaltimesnews/status/1468827585178001408

Huang ZA, Wang R (2019) Building a network to "tell China stories well'': Chinese diplomatic communication strategies on Twitter. Int J Commun 13:24

Hwang T, Pearce I, Nanis M (2012) Social bots: voices from the fronts. Interactions 19(2):38–45. https://doi.org/10.1145/2090150.2090161

Jacobs CS, Carley KM (2022) Taiwan: China's gray zone doctrine in Action. Small Wars J. https://smallwarsjournal.com/jrnl/art/taiwanchinas-gray-zone-doctrine-action

Kania E (2019) Chinese military innovation in the AI revolution. RUSI J 164(5–6):26–34. https://doi.org/10.1080/03071847.2019.1693803

King G, Pan J, Roberts M (2013) How censorship in China allows government criticism but silences collective expression. Am Polit Sci Rev 107(2):326–343

Krackhardt D, Stern R (1988) Informal networks and organizational crises: an experimental simulation. Soc Psychol Q 51(2):123–140. https://doi.org/10.2307/2786835

Luceri L, Deb A, Giordano S, Ferrara E (2019) Evolution of bot and human behavior during elections. First Monday. https://doi.org/10.5210/fm.v24i9.10213

McCombie S, Uhlmann A, Morrison S (2020) The US 2016 Presidential Election and Russia's troll farms. Intell Natl Secur 35(1):95–114. https://doi.org/10.1080/02684527.2019.1673940

Mozur P, Xiao M, Beltran G, Kao J (2021) China unleashed its propaganda machine on Peng Shuai's MeToo accusation. ProPublica. Accessed 16 Dec 2021

Mufumba I (2021) US sidesteps strategic ally Uganda on democracy summit. Daily Monitor. https://www.monitor.co.ug/uganda/news/national/us-sidesteps-stragetic-ally-uganda-on-democracy-summit-3633900?view=htmlamp. Accessed 5 Oct 2022

Myers SL, Mozur P, Kao J (2022) How bots and fake accounts push China's vision of winter Olympic wonderland. ProPublica. https://www.propublica.org/article/how-bots-and-fake-accounts-push-chinas-vision-of-winter-olympic-wonderland. Accessed 15 Dec 2021

Ng LH, Robertson D, Carley K (2022) Stabilizing a supervised bot detection algorithm: how much data is needed for consistent predictions? Online Soc Netw Media 28:100198. https://doi.org/10.1016/j.osnem.2022.100198

Phillips S, Uyheng J, Carley K (2022) Competing state and grassroots opposition influence in the 2021 Hong Kong Election. In: Thomson R, Dancy C, Pyke A (eds) Social, cultural, and behavioral modeling. SBP-BRiMS 2022. Lecture notes in computer science, vol 13558. Springer, Cham. https://doi.org/10.1007/978-3-031-17114-7_11

Ryan F, Bogle A, Zhang A, Wallis J (2021) #StopXinjiang rumors: the CCP's decentralised disinformation campaign. Australian Strategic Policy Institute. Policy Brief Report No. 54/2021

Schafer B (2021) China fires back at Biden with conspiracy theories about Maryland lab. Foreign Policy. https://foreignpolicy.com/2021/07/09/china-fires-back-at-biden-with-conspiracy-theories-about-maryland-lab/. Accessed 2 July 2022

Shu K, Sliva A, Wang S, Tang J, Liu H (2017) Fake news detection on social media: a data mining perspective. ACM SIGKDD Explor Newslett 19(1):22–36. https://doi.org/10.48550/arXiv.1708.01967

Spokesperson's Remarks (2021) Foreign Ministry Spokesperson's Statement on the "Summit for Democracy" Held by the United States. Ministry of Foreign Affairs of the People's Republic of China. https://www.mfa.gov.cn/mfa_eng/xwfw_665399/s2510_665401/202112/t20211211_10466939.html. Accessed 1 Oct 2022

Twitter Safety (2021) Disclosing state-linked information operations we've removed. https://blog.twitter.com/en_us/topics/company/2021/disclosing-statelinked-information-operations-we-ve-removed. Accessed 1 July 2022

Uren T, Thomas E, Wallis J (2019) Tweeting through the great firewall: preliminary analysis of PRC-linked information operations on the Hong Kong protests. Australian Strategic Policy Institute, Canberra

Waltman L, Van Eck N (2013) A smart local moving algorithm for large-scale modularity-based community detection. Eur Phys J B 86(11):1–14. https://doi.org/10.1140/epjb/e2013-40829-0

Woolley SC, Guilbeault D (2017) Computational propaganda in the United States of America: manufacturing consensus online. University of Oxford, Oxford

Woolley S, Howard P (2017) Computational propaganda worldwide: executive summary. Computational Propaganda Project. University of Oxford, Oxford

Yi C (2021) US-style #democracy creates more problems than solutions to the world. What give the US the confidence & makes it feel it's in a position to preach others on democracy? #NoMore #DemocracySummit #Whatisdemocracy #whodefinesdemocracy [Tweet]. Retrieved from https://twitter.com/CaoYi_MFA/status/1470006743199531012

Zhang L (2021) What if all the birds turn the same colour? Daily Monitor. https://www.monitor.co.ug/uganda/oped/commentary/what-if-all-the-birds-turn-the-same-colour--3645954. Accessed 15 Oct 2022

Charity S. Jacobs is a Ph.D. student at the School of Computer Science at Carnegie Mellon University and a Data Scientist for the U.S. Department of Defense. She holds a BA from the University of Wisconsin in political science and an MS from the University of Chicago in computational analysis and public policy. Her research interests include applied machine learning, natural language processing and disinformation detection.

Kathleen M. Carley Ph.D., is a professor of societal computing in the School of Computer Science at Carnegie Mellon University, an IEEE Fellow, the director of the Center for Computational Analysis of Social and Organizational Systems (CASOS), and the CEO of Netanomics. She is the 2011 winner of the Simmel Award from the International Network for Social Network Analysis and the 2018 winner of the National Geospatial-Intelligence Agency Academic Award from GEOINT.