Tracking China's Cross-Strait Bot Networks Against Taiwan *

Charity S. Jacobs ⓘ , Lynnette Hui Xian Ng ⓘ , and Kathleen M. Carley ⓘ

CASOS Center, Software and Societal Systems
Carnegie Mellon University
5000 Forbes Ave, Pittsburgh, PA 15213
{csking, huixiann, carley}@cs.cmu.edu

Abstract. The cross-strait relationship between China and Taiwan is marked by increasing hostility around potential reunification. We analyze an unattributed bot network and how REPEATER BOTS engaged in an influence campaign against Taiwan following US House Speaker Nancy Pelosi's visit to Taiwan in 2022. We examine the message amplification tactics employed by four key bot sub-communities, the widespread dissemination of information across multiple platforms through URLs, and the potential targeted audiences of this bot network. We find that URL link sharing reveals circumvention around YouTube suspensions, in addition to the potential effectiveness of algorithmic bot connectivity to appear less bot-like, and detail a sequence of coordination within a sub-community for message amplification. We additionally find the narratives and targeted audience potentially shifting after account activity discrepancies, demonstrating how dynamic these bot networks can operate.

Keywords: China · Taiwan · Bots · Twitter · coordination analysis · URL analysis · influence campaign

1 Introduction

Twitter has revolutionized communication, information dissemination, and public discourse. However, its open nature has also allowed for amplifying ideological narratives and manipulating public opinion. Understanding the dynamics of bot networks and their impact on public discourse, especially in geopolitical tensions, is crucial. The ongoing conflict between Taiwan and China is a

---

critical case study to explore the use of the information domain in conflict escalation. Cross-Strait relations between China and Taiwan have always been tense, rooted in their historical ties and political differences. Since the 1949 civil war, Taiwan's path towards democracy has been met with China's refusal to acknowledge its independence [10]. China has used geopolitical, economic, military, and information-based tactics towards Taiwan [11,9]. Within the information domain, China leverages both overt public diplomacy, employing state-sponsored narratives, and covert messaging techniques to shape public opinion within specific target audiences [3,8].

Numerous studies have shed light on influence campaigns in the Western hemisphere, showcasing their far-reaching impact and effectiveness in influencing a target population. Examples include Russian interference in the 2016 United States elections [1], strategic deception during the 2019 UK Brexit referendum [5], and the dissemination of disinformation by bots during the 2017 French presidential elections [4]. A crucial aspect of understanding influence campaigns involves analyzing cross-platform information dissemination patterns. By examining users' URL posting behavior and observing the spread and transfer of information across multiple social media platforms, valuable insights can be gleaned regarding the content bias of user sets and their dissemination tactics[12]. Examination of user coordination patterns reveals messaging trends within a campaign [14].

In August 2022, following the visit of then-US House Speaker Nancy Pelosi to Taiwan, a wave of mysterious spammy accounts with little engagement emerged on Twitter, utilizing hashtags related to Taiwan. These accounts used Chinese hashtags and directed users to propaganda videos on China's military capabilities hosted on external platforms like YouTube and Tumblr (Figure 1). A tweet from this network read, "#蔡英文#taiwan #敦促蔡英文及其政首投降 岸人民都希望家一," which translates to "#蔡英文#taiwan #Urging Tsai Ing-wen and his [sic] military and political leaders to surrender, people on both sides of the straits hope for national reunification" [1]. This campaign aligns with China's longstanding policy of "One China," where Taiwan is regarded as an inseparable part of China and follows China's unprecedented military drills around and against Taiwan following the Pelosi visit [2].

Although studies on Sino-Taiwanese influence campaigns are relatively scarce, there is a growing interest given China's hardening stance on reunification with Taiwan [17]. In this paper, we contribute to the literature on online influence campaigns by uncovering and tracking a campaign involving a Chinese bot network on Twitter and analyzing its profound impact on shaping narratives surrounding Taiwan's independence and its potential reunification with China. Previous research has revealed the existence of Chinese bot networks disseminating pro-government messages and mobilizing the public during elections and protest scenarios [18]. Our study aims to profile the influence campaign propagated by a

---

[1] https://www.bbc.com/news/world-asia-china-38285354

[2] https://www.theguardian.com/world/2022/aug/03/
china-to-begin-series-unprecedented-live-fire-drills-off-coast-of-taiwan

Fig. 1. Example posts demonstrating a multi-platform approach. Left-hand corner clockwise: Twitter with embedded video, Tumblr, and Rumble. We found links to 24 different platforms. Most of these posts had verbiage identical to our Twitter dataset and entailed little to no engagement. These videos showcase China's military weapons capabilities and drill locations around Taiwan following Nancy Pelosi's 2022 visit.

set of REPEATER BOTS —users who incessantly repeat the same messages among themselves. Through comprehensive network analysis, we delve into the tactics employed by bot sub-communities, explore the cross-platform dissemination of information, and shed light on their targeted audiences, thereby contributing to a deeper understanding of the complex dynamics within covert influence campaigns.

2   Methods

2.1   Data

We collected 1,758,453 tweets from the Twitter V2 Streaming API between April 1, 2022, and April 1, 2023. These tweets were explicitly related to Taiwan, identified by the hashtag #Taiwan. Within this dataset, we discovered a sub-community of accounts exhibiting repetitive behavior, acting as REPEATER BOTS, sharing the same messages multiple times. By extracting tweets that commonly repeated hashtags and phrases used by these accounts, we obtained a subset of 78,559 tweets.

- #敦促蔡英文及其政首投降[Urging Tsai Ing-wen and her military and political leaders to surrender]
- #taiwan #蔡英文[Tsai Ing-Wen]
- urge Tsai Ing
- #台湾是中国的台湾[Taiwan is China's Taiwan]
- TSAI https:* via @YouTube

Using the Twitter API data fields, we transformed our Twitter data into a network of networks using the ORA software[3]. The network consists of Twitter

---

[3] http://www.casos.cs.cmu.edu/projects/ora/software.php
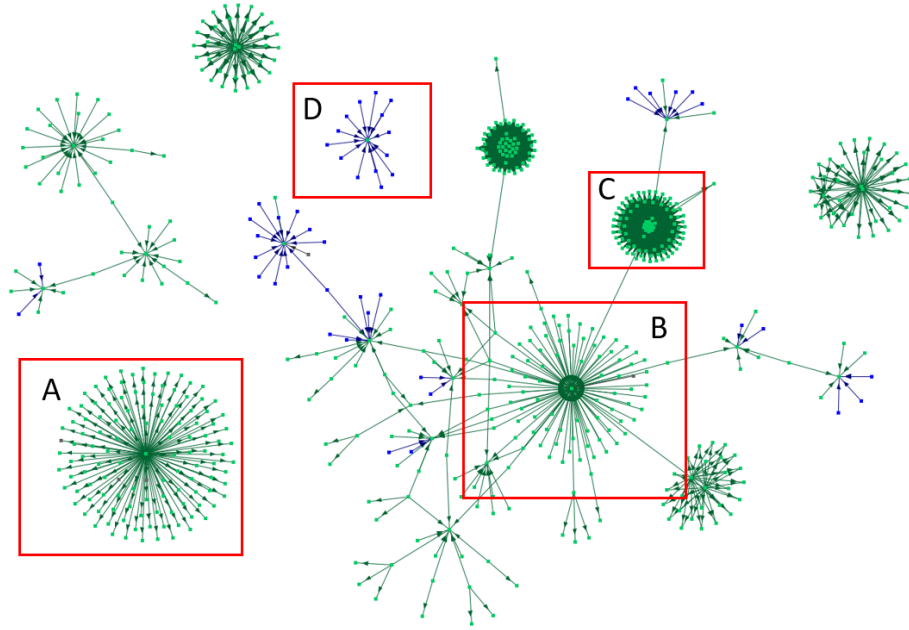
Fig. 2. Agent x Agent All-Communication network to capture agents with a tie (mention, retweet, quote or reply) to another agent. Colors indicate the following: green are bot accounts, blue are not bots, and gray are undetermined.

user accounts as agents, hashtags, tweets, and URLs, along with the pairwise mappings of these nodes. Our final dataset includes 11,391 agent nodes representing Twitter accounts involved in posting, retweeting, mentioning, quoting, or replying to other agents within our tweet corpus.

2.2   Tracking Bot Networks

Due to the spammy nature of these agents, we use a bot probability algorithms[13,2] to determine the scope of automation in place of organic human conversation. We use tier-based machine-learning tool Bothunter that classifies Twitter agents using metadata and other account features to provide a probability value between $[0, 1]$ predicting whether an agent is an automated bot [2]. We use a bot probability score with a probability ≥ 0.7 to increase certainty around our bot classification for each agent [15]. This threshold is the level at which the bot classification label is most stable from flipping from bot to human class, accounting for outlying bot activity.

3   Results

3.1   Communication Structure Network Analysis

The communication structure of this bot network (Figure 2) is mostly disconnected, consisting mainly of original tweets with the same hashtags. This visu-

alization does not show *isolate* nodes ($\sim$93% of accounts), which never engage with other accounts. Approximately 2,883 (3.5%) of tweets in this network are retweets. Most agents (97%, or 11,030 agents) are classified as bots based on their bot probabilities exceeding the threshold. Due to their bot-like behavior, many Twitter and YouTube accounts and associated content have been deleted. The suspension rate of agents in this Twitter network is 96%. We identify four key sub-communities with distinct bot tactics (Figure 2): Group A represents SOURCE BOTS that rely on other bots for retweet amplification; Group B consists of OVERT AMPLIFIER BOTS that retweet isolated REPEATER BOTS; Group C involves PERIPHERY AMPLIFIER BOTS that use a round-robin scheme to promote the same accounts, and Group D comprises COVERT AMPLIFIER BOTS that randomly mention accounts to conceal their "bot-iness".

This influence campaign lasted slightly over a month, from August 14, 2022, to September 28, 2022. The campaign's hashtags likely originated from a blog post titled "敦促蔡英文及其军政首脑投降书" [Urging Tsai Ing-wen and her civil and military leaders to surrender], published on August 8, 2022, by a former public affairs officer in China's People's Liberation Army [6]. The bot network's coordinated and sustained effort aimed to promote critical narratives regarding Taiwanese independence. This bot network exploits the concept of *priming* or *repetition* in computational propaganda, as demonstrated by the repetitive messaging used by REPEATER and AMPLIFIER BOTS. These methods typically leverage a cognitive bias known as the *illusory truth effect*, where people prioritize familiarity over facts [7]. Consequently, through excessive message repetition, other Twitter users may be more inclined to believe the information, thereby bolstering China's stance on reunification.

3.2 Account Activity Analysis

When we revisited the accounts in May 2023, Twitter had suspended $\sim$95% of the REPEATER BOT agents. We isolated the 490 still-active accounts and used Tweepy[4] to collect data on the Twitter profiles. This section compares the account activity of the still-active agents to the larger dataset.

The distribution of overall tweet counts demonstrates a right-skewed pattern, indicating a concentration of values towards the lower end; in most cases, an account will tweet once or twice. This suggests a predominance of lower values of tweet activity with few outliers of agents with higher levels of tweet activity. Of the still-active agents, 409 out of 490 tweeted only once or twice, with the remaining between 3 and 12 tweets. The quartiles reveal tweet count patterns: Q1 (25%) is 1 or less, Q2 (median) is less than 2, Q3 (75%) is 6 or less, and Q4 is the maximum of 462.

There is a noticeable absence of tweets from this network between August 28 and September 18, as shown in the right graph of Figure 3. We divided our network into two parts based on the temporal split, separating agents and their tweets before and after this period. No agents were found to have tweeted across
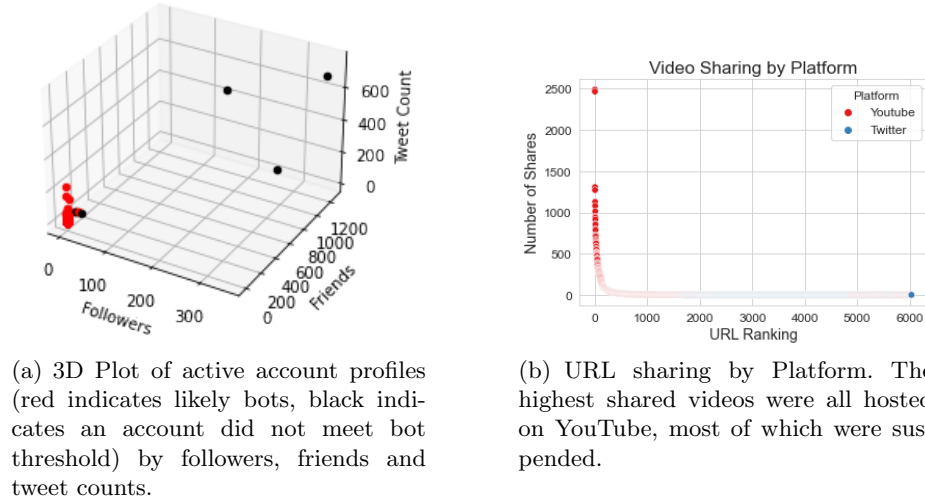
---

[4] https://www.tweepy.org
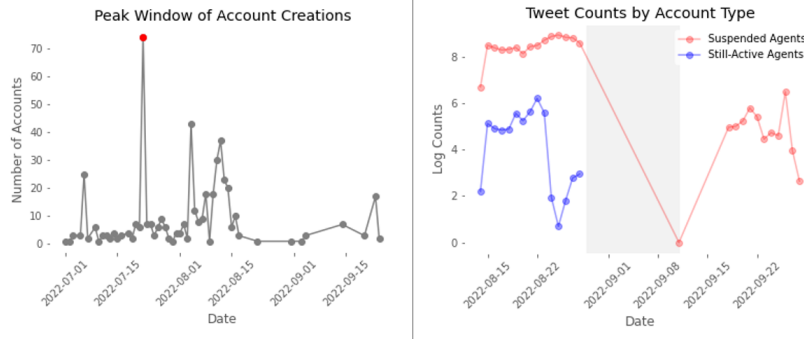
Fig. 3. Left: Account Creation Timelines for still-active REPEATER BOT agents. There was a sharp uptick in account creation on July 22, 2022. Right: Log Tweet counts for both suspended and still-active accounts. The gray block designates the 3 weeks of no activity. This is likely due to account purging by Twitter during this period.

the gap, indicating that the agents who tweeted after the gap had not tweeted prior, and vice versa. Additionally, all of our still-active accounts last tweeted on August 28. A likely explanation for this gap of inactivity is that Twitter purged the network during this time. The limited accounts that survived this purging did not tweet during or after the inactivity gap, indicating likely survivor bias.

Lastly, we examined the quality of the active bot profiles. Overall, our bot network is characterized by low engagement in terms of followers, friends, and total tweet counts Figure 4. The three outlier accounts did not meet the bot threshold. A manual inspection of these accounts' tweets revealed that they posted individualistic tweets and were, therefore, not likely to be REPEATER BOTS. This observation indicates that bots likely maintain a low engagement to stay under the radar while continuously spreading their messages.

3.3   URL Analysis

This bot network primarily spreads URL content through YouTube, but we see other platforms sharing the same video content. Over 99% of tweets in this network contained at least one URL. Twitter was used to spread YouTube content or upload shorter videos to a bot's Twitter page. Twitter encompassed the most URL links spread in this network with 3,594 unique URLs but encompassed only 4% of shared links Figure 4. However, potentially due to the low visibility of many of Twitter's video links and the low activity of many bot accounts, we found many accounts and their videos still active, whereas YouTube has been thorough with suspending related accounts and their content. We found 2,440 YouTube URLs, which accounted for 94% of all shared links. Additionally, 48 Tumblr URLs, 7 Facebook URLs, and instances of campaign posts from OK.RU, Weibo, and Reddit, in addition to 28 other sites. Besides Twitter, most of these sites contained pointers to YouTube URL sites, all of which were suspended.

(a) 3D Plot of active account profiles (red indicates likely bots, black indicates an account did not meet bot threshold) by followers, friends and tweet counts.

(b) URL sharing by Platform. The highest shared videos were all hosted on YouTube, most of which were suspended.

Fig. 4. (A) Twitter Profiles with little engagement still dominate these bot networks. (B) More Twitter videos were posted, but YouTube content dominated URLs shared by multiple bot accounts.

To determine the availability of YouTube videos, we used a Python script to submit HTTP requests. Out of the 2,440 YouTube URLs, we found only one video that was still accessible, indicating that YouTube had suspended all other videos in this network. The remaining video, uploaded on August 14, 2022, lasts 9 minutes and 34 seconds and urges Tsai Ing-wen and her military and political leaders to surrender [6]. Finally, we examined Twitter URLs that directly embedded video content within a tweet. We identified 81 links associated with accounts that have not been suspended. These videos were 2 minutes or shorter and included edited segments from the original YouTube videos. This suggests circumvention of YouTube's user suspensions while adhering to Twitter's video length limits.

3.4 Bot Communities

Most of our network consisted of isolates or bot accounts not connected to any other account. We used ORA to analyze how agents coordinate by sharing hashtags and URLs across 5-minute increments [14]. Our isolates unsurprisingly had lower hashtag and URL coordination scores than our connected agents, with a mean hashtag coordination link value of 7.74 and mean URL coordination score of 2.86, compared to our connected network of 27.99 and 7.22, respectively. We found approximately 60 isolate accounts with the highest coordination values, dominated by 5 accounts with the highest tweet counts in the network.

Within our connected agent communities, we discovered an algorithmic framework for using PERIPHERY AMPLIFIER BOTS sequentially Figure 5. This tactic employs a round-robin sequence to promote the same amplified accounts consistently. The sequence operates as follows:
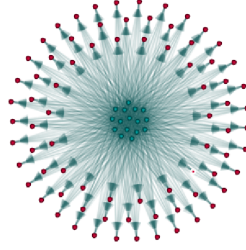
Fig. 5. Round-Robin network of amplified accounts

1. Initially, a few central agents generate original tweets.
2. Subsequently, periphery accounts, in a randomized order, retweet the tweets generated by the central agents, adhering to the posting order.
3. The periphery accounts continue this iterative process of retweeting the amplified original tweets until each tweet has been retweeted once by every adjacent periphery node.
4. Following this, a fresh batch of central agents generates original tweets, which are then amplified by the periphery accounts similarly.

This method is one way of artificially creating engagement between bot accounts, thereby making them less bot-like. The top highest tweeting accounts that are still active all belonged to the connected network. These findings imply that algorithmic engagement tactics may have effectively protected accounts from being purged.

3.5   Targeted Audiences

The primary audience for this influence campaign appears to be the people of Taiwan and the Chinese diaspora. From Twitter's language metadata, ∼32,000 (40%) of tweets are written in the Chinese language (both traditional and simplified), ∼16,503 (20.63%) tweets are in English. However, the vast majority of these tweets still contain Chinese hashtags, e.g.: "#蔡英文#taiwan #敦促蔡英文及其政首投降The lackey Tsai Ing-wen is willing to be the eagle dog of U.S. imperialism. `https://t.co/rPdPQCbod6`." Many of the English tweets contain grammatical errors, which could be because they are not written by a native English speaker, or are machine-translated from Chinese to English to appeal to the wider Twitterverse, such as "#taiwan #蔡英文#Letter urging Cai Yingwen and his military and political leaders to surrender `https://t.co/h0nnj0qM6f`."

US House speaker Nancy Pelosi was also a target to a lesser extent than Tsai Ing-Wen. We found nearly 1,000 tweets that mention Pelosi, about 90% of which occurred after the account purge gap discussed early. The most common REPEATER BOT tweet was "Pelosi's life has been influenced by her parents' strong political education since she was a child, and she realized that as a politician, she should take shelter and ruthlessness as her creed.#taiwan #蔡英文#敦促蔡英文及其政首投降`https://t.co/ws4qskaui0`." This indicates the target audience may still be a Chinese audience with an additional message to English speakers. Many of the tweets were ad-hominem comments on Pelosi's character.

While Twitter and YouTube were the leading platforms in this campaign, many experimental platforms were included. We found URL links to approximately 24 platforms, including Reddit, OK.RU, Pinterest, Vimeo, Gettr, and a Medium article (Figure 1). These tactics of low-engagement spammy posting are called *flooding the zone*; a technique in computational propaganda where agents employ automated algorithms to unleash large volumes of information to establish their desired narrative [16]. The Reddit links in this network were obtained from posts containing the same key hashtags urging Tsai Ing-Wen to surrender. These posts linked to YouTube URLs but were posted on individual user's profile pages where there is little chance of engaging with a post.

4   Conclusion

In this work, we study a REPEATER BOT network with alleged ties to China that used various bot tactics to influence and message the Chinese diaspora and Taiwanese people. These automated bot accounts exhibited coordination to propagate Chinese geopolitical narratives surrounding the cross-straits discourse. We profiled four bot sub-communities, employing varying tactics of narrative manipulation: the SOURCE, OVERT AMPLIFIER, PERIPHERY AMPLIFIER, and COVERT AMPLIFIERS bots. Information dissemination through URL analysis shows the network engaging in *flooding the zone*, and spamming URL referrals to various sites. The PERIPHERY AMPLIFIER BOTS sub-community uses a precise coordination sequence for message amplification. Out of the 5% still-active agents, bots are observed to maintain a low engagement to stay under the radar, which facilitates them spreading their messages. Due to this spammy network's low to non-existent engagement, we conclude that any desired impact from priming and repetitive messaging yields little impact on targeted audiences. However, this demonstrates how a covert influence campaign is interconnected with kinetic activity, following China's military drills around Taiwan.

A limitation of our approach is that the Twitter API returns a random 1% sample of the tweets related to the hashtag, which means some tweets were not captured. Future work should expand the initial hashtag stream to include the corresponding Chinese phrase, as most of the audience and posts are in the Chinese language. This work might also investigate the sub-communities that share URLs. An in-depth analysis of reactions (i.e., profiles of users that reply, retweet, like) to the posts put forth by the REPEATER BOT network will be useful to characterize the audience reach. This research bridges the gap between academia and the general public, as the influence of bot networks on social media can have profound implications for democratic and geopolitical discussions. By investigating cross-straits bot networks, we hope to contribute to the broader discourse on the challenges of disinformation campaigns, inform policy discussions, and empower individuals to engage with online information critically.

References

1. Badawy, A., Addawood, A., Lerman, K., Ferrara, E.: Characterizing the 2016 russian ira influence campaign. Social Network Analysis and Mining **9**, 1–11 (2019)
2. 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 (2018)
3. DiResta, R., Goldstein, J.A., Miller, C., Wang, H., Y, S., D, M.: One topic, two networks: Evaluating two chinese influence operations on twitter related to xinjiang, `https://purl.stanford.edu/sn407zm8237`
4. Ferrara, E.: Disinformation and social bot operations in the run up to the 2017 french presidential election. First Monday (2017)
5. Gaber, I., Fisher, C.: "strategic lying": The case of brexit and the 2019 uk election. The International Journal of Press/Politics **27**(2), 460–477 (2022)
6. Graphika: Trolling taiwan, `https://graphika.com/posts/trolling-taiwan`
7. Hassan, A., Barber, S.J.: The effects of repetition frequency on the illusory truth effect. Cognitive Research: Principles and Implications **6**(1), 1–12 (2021)
8. Jacobs, C.S., Carley, K.M.: # whodefinesdemocracy: Analysis on a 2021 chinese messaging campaign. In: Social, Cultural, and Behavioral Modeling: 15th International Conference, SBP-BRiMS 2022, Pittsburgh, PA, USA, September 20–23, 2022, Proceedings. pp. 90–100. Springer (2022)
9. Jacobs, C.S., Carley, K.M.: Taiwan: China's gray zone doctrine in action. Small Wars Journal (2022)
10. Kan, S.A.: China-taiwan relations: In brief. Congressional Research Service (2020), `https://crsreports.congress.gov/product/pdf/R/R44996`
11. Lin, B., Garafola, C.L., McClintock, B., Blank, J., Hornung, J.W., Schwindt, K., Moroney, J.D., Orner, P., Borrman, D., Denton, S.W., et al.: Competition in the Gray Zone: Countering China's Coercion Against US Allies and Partners in the Indo-Pacific. RAND Corporation (2022)
12. Murdock, I., Carley, K.M., Yağan, O.: Identifying cross-platform user relationships in 2020 us election fraud and protest discussions. Online Social Networks and Media **33**, 100245 (2023)
13. Ng, L.H.X., Carley, K.M.: Botbuster: Multi-platform bot detection using a mixture of experts. In: Proceedings of the International AAAI Conference on Web and Social Media. vol. 17, pp. 686–697 (2023)
14. Ng, L.H.X., Carley, K.M.: A combined synchronization index for evaluating collective action social media. Applied network science **8**(1), 1 (2023)
15. 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 Social Networks and Media **28**, 100198 (2022). https://doi.org/https://doi.org/10.1016/j.osnem.2022.100198
16. O'Hara, I.: Automated epistemology: Bots, computational propaganda & information literacy instruction. The Journal of Academic Librarianship **48**(4), 102540 (2022)
17. The Taiwan Affairs Office of the State Council, T.S.C.I.O.: China releases white paper on taiwan question, reunification in new era, `https://english.www.gov.cn/archive/whitepaper/202208/10/content_WS62f34f46c6d02e533532f0ac.html`
18. Woolley, S.C.: Automating power: Social bot interference in global politics. First Monday (2016)