

Chinese computational propaganda: automation, algorithms and the manipulation of information about Chinese politics on Twitter and Weibo

Gillian Bolsover and Philip Howard¹

Oxford Internet Institute, University of Oxford, Oxford, UK.

gillianbolsover@gmail.com, philip.howard@oii.ox.ac.uk

1 St Giles, Oxford, OX1 3JS, United Kingdom

+44 (0) 7503 646092

Gillian Bolsover is a research associate at the University of Oxford's Oxford Internet Institute. She completed her DPhil at the Internet Institute in 2017 and holds a dual MSc/MA in Global Media and Communications from the London School of Economics and Political Science and Fudan University in Shanghai, China.

Philip Howard is a Professor of Internet Studies and the Director of Research at the Oxford Internet Institute. He is the Principal Investigator of the Computational Propaganda research project that investigates how bots, algorithms and other forms of automation are used by political actors in countries around the world.

¹ The authors gratefully acknowledge the support of the European Research Council, "Computational Propaganda: Investigating the Impact of Algorithms and Bots on Political Discourse in Europe," Proposal 648311, 2015-2020, Philip N. Howard, Principal Investigator. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the European Research Council.

Chinese computational propaganda: automation, algorithms and the manipulation of information about Chinese politics on Twitter and Weibo

A 2016 review of literature about automation, algorithms and politics identified China as the foremost area in which further research was needed because of the size of its population, the potential for Chinese algorithmic manipulation in the politics of other countries, and the frequency of exportation of Chinese software and hardware. This paper contributes to the small body of knowledge on the first point (domestic automation and opinion manipulation) and presents the first piece of research into the second (international automation and opinion manipulation). Findings are based on an analysis of 1.5 million comments on official political information posts on Weibo and 1.1 million posts using hashtags associated with China and Chinese politics on Twitter. In line with previous research, little evidence of automation was found on Weibo. In contrast, a large amount of automaton was found on Twitter. However, contrary to expectations and previous news reports, no evidence was found of pro-Chinese state automaton on Twitter. Automation on Twitter was associated with anti-Chinese state perspectives and published in simplified Mandarin, presumably aimed at diasporic Chinese and mainland users who ‘jump the wall’ to access blocked platforms. These users come to Twitter seeking more diverse information and an online public sphere but instead they find an information environment in which a small number of anti-Chinese state voices are attempting to use automation to dominate discourse. Our understanding of public conversation on Twitter in Mandarin is extremely limited and, thus, this paper advances the understanding of political communication on social media.

Keywords: Twitter; Weibo; China; bots; politics; computational propaganda

Introduction: The rise of computational propaganda and social media bots

Twenty-sixteen has come to be seen as a time of political turmoil and the year in which long-standing fears about the negative effects of social media on democratic politics were finally realised. In a referendum marred by false promises based on misleading information (Helm, 2016), growing nationalism that led to the murder of an MP (Cobain & Taylor, 2016) and the algorithmic manipulation of online public opinion (Howard & Kollanyi, 2016), the UK narrowly voted to leave the EU.

Several months later, polemical billionaire Donald Trump won the US presidency. During campaigning, automated accounts, particularly in pro-Trump hashtags, dominated discourse on Twitter (Howard, Kollanyi, & Woolley, 2016) and junk news was shared as frequently as professionally-produced news (Howard, Bolsover, Kollanyi, Bradshaw, & Neudert, 2017). Accusations of Russian technological interference in the election are now the subject of several major congressional investigations (LoBianco, 2017).

Although the true influence of automated (bot) accounts on social media is unknown, emerging evidence suggests that they are effective at spreading information and deceiving users. In the run-up to the US Presidential election, human Twitter users retweeted bots at the same rate as other humans (Bessi & Ferrara, 2016). It has also been shown that typical Internet users cannot determine whether information has been produced by a human or a bot (Everett, Nurse, & Erola, 2016).

Although bots were identified in US political events as early as 2010 (Mustafaraj & Metaxas, 2010; Ratkiewicz et al., 2011), the need to understand bots and their effects is now more urgent. Technical and policy solutions to the apparent problem of bots have been advancing ahead of academic research and there are several notable areas in which knowledge is lacking. Chief among these is understanding computational propaganda in relation to China, which was identified as the primary area in need of further

investigation in a review of literature concerning automation, algorithms and politics (Shorey & Howard, 2016).

Media reports of Chinese computational propaganda

As yet, no academic research has investigated whether the Chinese state uses bots as part of its international propaganda strategy. However, there have been sporadic media reports of Chinese state-associated bot activity and some academic reviews of media reports concerning social media manipulation.

A 2016 review of 48 English-language newspaper reports concluded that in authoritarian countries bots tend to be used to demobilise opposition voices and spread pro-government messages, while in countries with a longer history of democracy they are generally only used for social media follower padding (Woolley, 2016). A similar review of 83 English-language media reports concluded that authoritarian states tend to focus on their domestic populations, while democratic countries frequently use social media manipulation to target foreign publics (Bradshaw & Howard, 2017).

However, this conclusion (based on a limited number of English-language media reports) that authoritarian countries do not use automation to target foreign populations contrasts with the current concern about Russian computational propaganda. A US Intelligence report concluded that Vladimir Putin targeting the 2016 US Presidential Election with a multifaceted influence campaign that blended “covert intelligence operations—such as cyber activity—with overt efforts by Russian Government agencies, state-funded media, third-party intermediaries, and paid social media users or ‘trolls’” (Intelligence Community Assessment, 2017, p. 2).

Little scholarly attention has been paid to whether China undertakes similar media manipulation strategies. However, media reports have suggested that the Chinese state may be attempting to influence public opinion on Twitter. In early 2014, it was reported that more than 100 fake Twitter accounts were spreading positive propaganda in

English about conditions in Tibet; these accounts were followed by many human users, who apparently believed these accounts belonged to real people (Kaiman, 2014).

Later that year, there was an alleged bot attack on the actor Murong Xuecun, who had been critical of the Chinese state; more than 800 recently created Twitter accounts circulated a 10-page article attacking the actor (Henochowicz, 2014; Phillips, 2014). A similar incident was reported in October 2017, when numerous apparently automated accounts posted messages attacking the Chinese businessman and anti-corruption campaigner Guo Wengui (Collins & Cox, 2017).

These media reports suggest that China may be using automation to spread propaganda but no academic work has investigated this issue. However, the body of academic work on China's foreign media strategy more broadly may be relevant to understanding whether the state might use bots and automation to spread propaganda.

Chinese soft power, public diplomacy and foreign propaganda

In the early 2000s, China intensified its focus on its foreign image and started to cultivate consent for the country's peaceful rise, using official state media to engage with civil society in foreign countries (Y. Wang, 2008). The 2006 Five-Year Plan argued China's soft power should be based on "strong propaganda methods and strong propaganda capabilities" (Hayden, 2012, p. 137).

However, this propaganda has focused on traditional media, paying little attention to online media (Creemers, 2015). Between 2009-2010, the Chinese government reportedly spent \$8.7 billion on foreign propaganda, with the majority going to China Central Television, China Radio International, the Xinhua News Agency and the China Daily newspaper (Shambaugh, 2010).

While these big four providers are common names, there is also evidence of covert strategies. A 2015 Reuters investigation uncovered 33 radio stations in 14

countries broadcasting pro-Chinese state propaganda and structured so as to obscure that the majority shareholder was China Radio International (Qing & Shiffman, 2015).

The majority of the academic work on Chinese foreign propaganda points to a focus on traditional media. However, conditions change rapidly in China. Xi Jinping, who took over the helm of the party in late 2012, has taken a hard-line attitude towards domestic media liberalisation and this appears mirrored in foreign propaganda efforts.

Between the time Xi took office and December 2015, the Freedom House noted more than 40 instances in 17 countries and international institutions of Chinese information controls negatively affecting free expression outside China (Cook, 2015). There have also been reports of interference in Chinese language media in countries such as Canada and Australia (Kalathil, 2017).

Xi's crackdown on Chinese online information combined with the rising prominence of the Internet suggests that the online might have become a greater part of China's external media strategy. In the lead-up to China's 2016-2020 Five Year Plan, the concept of Internet Power was prominent in guideline documents (Livingstone, 2016). It also seems that Chinese production of online propaganda, such as Internet memes, clickbait headlines and promotional videos, has increased (Livingstone, 2016; Chow, 2017). These media are instances of computational propaganda and suggest the Chinese government is paying more attention to foreign social media; however, there has been no academic research to investigate whether the bots and automation that were so prominent in recent political events in the US are being employed to disseminate Chinese foreign propaganda.

Domestic propaganda and opinion manipulation in China

China has a long history of information control and a very different approach to propaganda. Since the communist revolution, the media have been run on a Marxist model that puts the needs of the state above truth, impartiality or diversity (Li, 2013;

Xinhua, 2016). After the rise of the Internet, these ideas were first extended to social media companies, then online opinion leaders and finally all Internet users (Bolsover, 2017).

Many of the techniques used to control content on the Chinese Internet are automated (Ng, 2015; Zhu, Phipps, Pridgen, Crandall, & Wallach, 2013). However, little evidence exists for the bots that have been prominent in other countries. For years, commentators spoke about the ‘50-Cent Party,’ individuals paid 50 cents per post to attack critics and support the state online (Greitens, 2013; Hassid, 2012).

However, based on a leak from an Internet Propaganda Office, a research team at Harvard came to a surprising conclusion; rather than an army of users paid by the post, the 50-Cent Party was composed of government employees who posted pro-state content as part of their regular jobs (King, Pan, & Roberts, 2017). Investigating whether these posts were automated, the team concluded “the evidence strongly indicates to the contrary” (ibid, p. 11).

Despite a lack of evidence of automation, fake accounts appear to be frequently employed to manipulate information on the Chinese microblogging giant Sina Weibo. An analysis of networks of news dissemination found that retweeting by fake accounts occurred in 6% of news stories and that 30% of the accounts that acted as opinion leaders were fake (Bolsover, 2013).

Although fake accounts are frequently employed to manipulate public opinion, there has been no evidence of automation in China. This conclusion is somewhat surprising given the sophistication of Chinese Internet control and the prevalence of use of bots in other countries. Although the Harvard study found no evidence of automation, it was based on a single leak from one local-level Internet propaganda office. Thus, more research is necessary to establish whether or not there is bot activity on Chinese domestic social media.

Methods and data collection

Social media are the most widely used functionality of the contemporary Internet. Of social media platforms, microblogs are an ideal venue for the investigation of online computational propaganda because of their public nature. Almost all of the previous research about bots and automation has focused on Twitter. Thus, this research focuses on Twitter and its domestic counterpart in China, Sina Weibo.

Researching Computational Propaganda on Weibo

Although sometimes referred to as a Chinese Twitter, Sina Weibo², the largest microblogging platform in China, provides different technical and social affordances for political speech and public opinion manipulation. A particular affordance of Weibo that does not have a parallel on Twitter is its threaded commenting system, which provides a space for users to engage in discussions that are more akin to those that occur on Facebook (Bolsover, 2016). A quarter of all “50-Cent Party” posts made in Weibo comments (King et al., 2017). Thus, Weibo comments are a prime venue in which automated computational propaganda might occur.

In order to investigate whether evidence of computational propaganda appears in Weibo comments, the posts of 26 major information providers—news organizations, government departments and official mouthpieces—were collected over the 2017 Spring Festival period. These accounts were selected to cover the largest state providers of news information on the platform, drawing from the platform’s leader boards and lists of the highest circulation media providers in China. Prior research has suggested that there are higher levels of state-led public opinion manipulation during official holidays (King et al., 2017).

² Weibo literally means microblog and several commercial microblogging platforms exist, including those of Sina and Tencent. However, Sina Weibo is the largest microblogging platform in China and is often simply referred to as Weibo. In line with this discourse, further references in this paper to Weibo (capitalised) should be understood as referring to the Sina Weibo platform.

Table 1: The 26 selected information providers and their reach

Account name	English name	Number of followers ³
人民日报	People's Daily	55.7 million
头条新闻	Weibo breaking news channel	52.6 million
央视新闻	CCTV News	52.3 million
人民网	People.cn	39.8 million
新华网	Xinhua	31.4 million
新华视点	Xinhua Viewpoint	30.5 million
中国日报	China Daily	30.1 million
公安部打四黑除四害	Ministry of Public Security and Public Security Bureau targeting counterfeit, fake and stolen goods and gambling and drug-related crimes	29.2 million
新浪娱乐	Weibo entertainment channel	22.7 million
光明日报	Guangming Daily	19.0 million
微天下	Weibo 24-hour Information Channel	16.2 million
新浪财经	Weibo economics channel	14.9 million
新浪科技	Weibo Science and Technology Channel	12.0 million
南方日报	Southern Daily	11.2 million
环球时报	Global Times	9.0 million
新浪视频	Weibo video channel	8.7 million
北京青年报	Beijing Youth Daily	8.0 million
江宁公安在线	Nanjing Public Security Bureau, Jiangning Branch	8.0 million
广州公安	Guangzhou Province Public Security Bureau	5.7 million
央视网	CCTV	5.2 million
中国网络电视台	Chinese Network Television	3.7 million
新疆发布	Xinjiang Propaganda Department	3.7 million
凤凰网	Phoenix News	2.8 million
红旗文稿	Red Flag Manuscripts	610,000
青春上海	Shanghai Youth League	413,000
拉萨发布	Lhasa, Tibet Propaganda Department	201,000

All of the posts made by these 26 information providers were collected between 26 January and 7 February 2017 (n=6,145). Comment data for each of these posts was collected at least two weeks after they were originally posted. The final dataset contained 1,543,165 comments by 815,776 unique users.

Researching Computational Propaganda on Twitter

Although Twitter is blocked in China, it is still used by some Chinese individuals, particularly as a subversive space for those who want to engage in discussion about sensitive issues (Sullivan, 2012). Geolocation of a random sample of Twitter accounts found that about 0.17% of all monthly-active users were located in mainland China

³ As of January 2018.

(Bolsover, In Press). Furthermore, as described in previous sections, the Chinese state actively cultivates a positive image of the country among foreign populations and there have been several media reports of bot activity associated with Chinese state interests on Twitter.

Thus, in order to investigate Chinese computational propaganda on Twitter, a preliminary list of hashtags associated with China and Chinese politics was drawn up. All of the tweets made between 24 January and 5 February 2017 using one of these hashtags was collected. These tweets and their concurrent hashtags were analyzed to ascertain hashtags commonly used to post about Chinese politics. A final list of 27 of the most common hashtags associated with Chinese social, political and cultural issues was established (Table 2). All of the tweets posted between 21 February and 8 April 2017 that used one of these hashtags was collected.

Table 2: The hashtags used for data collection on Twitter

Hashtags Collected	Description
#China, #Hongkong, #Beijing, #Shanghai, #Xinjiang, #Tibet, #Taiwan	Important locations (English)
#中国, #香港, #北京, #上海, #新疆 #西藏 (China, Hong Kong, Beijing, Shanghai, Xinjiang and Tibet)	Important locations (Mandarin)
#ChinaCulture, #ChinaTravel, #panda	Positive foreign publicity
#SouthChinaSea, #Diaoyudao, #Senkaku	Areas of territorial disagreement
#dalailama, #buddhism, #Kadampa	Buddhism
#XiJinping, #习近平, #XiVisit	Chinese premier Xi Jinping
#人权 (Human rights)	
#AntiChina	

Computational propaganda on Twitter: a dominance of anti-state voices

The final dataset contained 1,177,758 tweets from 254,132 unique accounts. Quantitative analysis using custom Python scripts revealed that information about China and Chinese politics on Twitter is dominated by a small number of voices. More than half of the tweets were made by users who posted more than 100 times during the data collection period and 42% of posts were posted by users who posted more than 300 times. Almost 30% of the tweets in the dataset came from the top 100 highest-posting users.

Data returned from the Twitter (as well as the Weibo) API provides the source platform of the tweet, such as Twitter for iPhone, the Twitter web client, or third-party platforms such as TweetDeck or Hootsuite. These data can provide the best evidence for account automation; if 100% of the account's tweets are made using an automation platform it is, without a doubt, a bot. Seventy-one of the top-100 highest posting accounts posted all or almost all of their posts using known automation platforms: 35 used the Japanese platform twittbot.net, nine IFTTT (If This Then That) and four dlvr.it. Additionally, many of these accounts appeared to be using custom automation scripts.

This provides a clear indication that there is significant automation within this dataset. However, because automation can be executed through custom scripts or via a standard client such as Twitter for Android or iPhone, using only post source to identify bots, particularly if this process is automated, will likely produce false negatives. Thus, in order to further investigate evidence for automation in the dataset and to evaluate the effectiveness of quantitative, scalable methods for identifying bots, two metrics used in previous research were applied to the dataset.

The tool BotOrNot (now Botometer) was developed by researchers at Indiana University. A score of 50% or higher on BotOrNot is generally seen as indicating the account is "suspicious to a scrupulous analysis" (Bessi & Ferrara, 2016). The average BotOrNot score of these 100 accounts was 54.7%, indicating a relatively high level of bot activity. Twenty-two of the top 100 posting accounts had a BotOrNot score of less than 50; however, these accounts clustered at the upper end of the range with seven accounts scoring 48 or 49. However, several of the accounts that scored less than 50 were clearly bots, with 100% of their tweets posted using automation platforms.

Another quantitative, scalable measure that has been used to identify automated accounts is posting frequency; a cut-off point of 50 posts per day in monitored hashtags was used to identify likely automated accounts in the 2016 US election (Howard et al.,

2016). The top 100 highest posting users in the Twitter dataset posted on average 70 times per day, with the top 38 highest posting users posting more than 100 times per day. However, many accounts posting only through automation platforms or that received high BotOrNot scores, posted less than 50 times per day across the examined hashtags.

Each of these three metrics – post source, BotOrNot and post frequency – suggests high levels of automation among the highest posting users, who produced almost 30% of the posts in the dataset. The comparison of the three metrics suggest that each is conservative. They are unlikely to produce false positives but may produce false negatives. Post source is the most reliable method for bot identification but it is not scalable over large datasets.

A further limitation of these methods is that they focus solely on quantitative data. This can help identify bots and the hashtags in which they are active but cannot speak to the actual content that these bots are associated with, i.e. the propaganda they might posts and the interests furthered by this automation. It is important to remember that not all bots promote propaganda. Institutions, companies, news media and individuals all use automation to post non-propaganda content. Thus, in order to understand the nature of computational propaganda about China on Twitter, it is necessary to qualitatively analyze the profiles and posts of these high-posting accounts.

Previous research has found evidence of likely automation based on numerous characteristics: posting frequency (bots tend to post much more frequently than individual users), post time (bots can post consistently across the entire day while humans need to sleep), post content (bots often post only about a single issue), post repetitiveness (bots often repeatedly post the same or similar messages), percentage of retweets (many bots only retweet other's content), connectivity (bots are often part of groups that interact with each other through mutual following and sometimes retweeting), number of friends and followers (many bots build followers through reciprocal

relationships and thus have a similar number of friends and followers, other bots will have almost no friends or followers) and post interaction (many bots will have no user interaction on their timelines).

The profiles of each of the top-100 highest posting users was inspected and evaluated according to the above metrics. Based on this examination, each of the 100 accounts that had not been suspended by the time of analysis ($n = 82$) was deemed to be an automated account⁴. The type of content posted by these accounts was coded according to a scheme derived from an examination of the dataset. No accounts posting pro-Chinese-state content were found within these 100 users; however, half of these accounts posted anti-Chinese-state content. Among these there were two large groups: the 1989 group and the pan-Asia group (Table 3). This is a surprising finding given previous media reports of Chinese state bot activity on Twitter and, thus, descriptions of each of these two groups are provided in the following sections.

The 1989 bot group

Accounts in this group promote content about human rights in China, particularly related to keeping alive the memory of the 1989 student-led democracy movement that ended with the Tiananmen Square “incident”. All of the posts of accounts in this group are in simplified Chinese and information posted by these accounts dominates hashtags related to China and major Chinese cities in both English and simplified Mandarin (#China, #Hongkong, #Beijing, #Shanghai, #香港, #北京, #上海).

⁴ The fact that 18 of the accounts had been deleted between data collection and the qualitative analysis phase suggests that these accounts, which were predominantly automated using custom scripts, were identified as bots and deleted by the platform.

Table 3. Top 100 highest-posting accounts

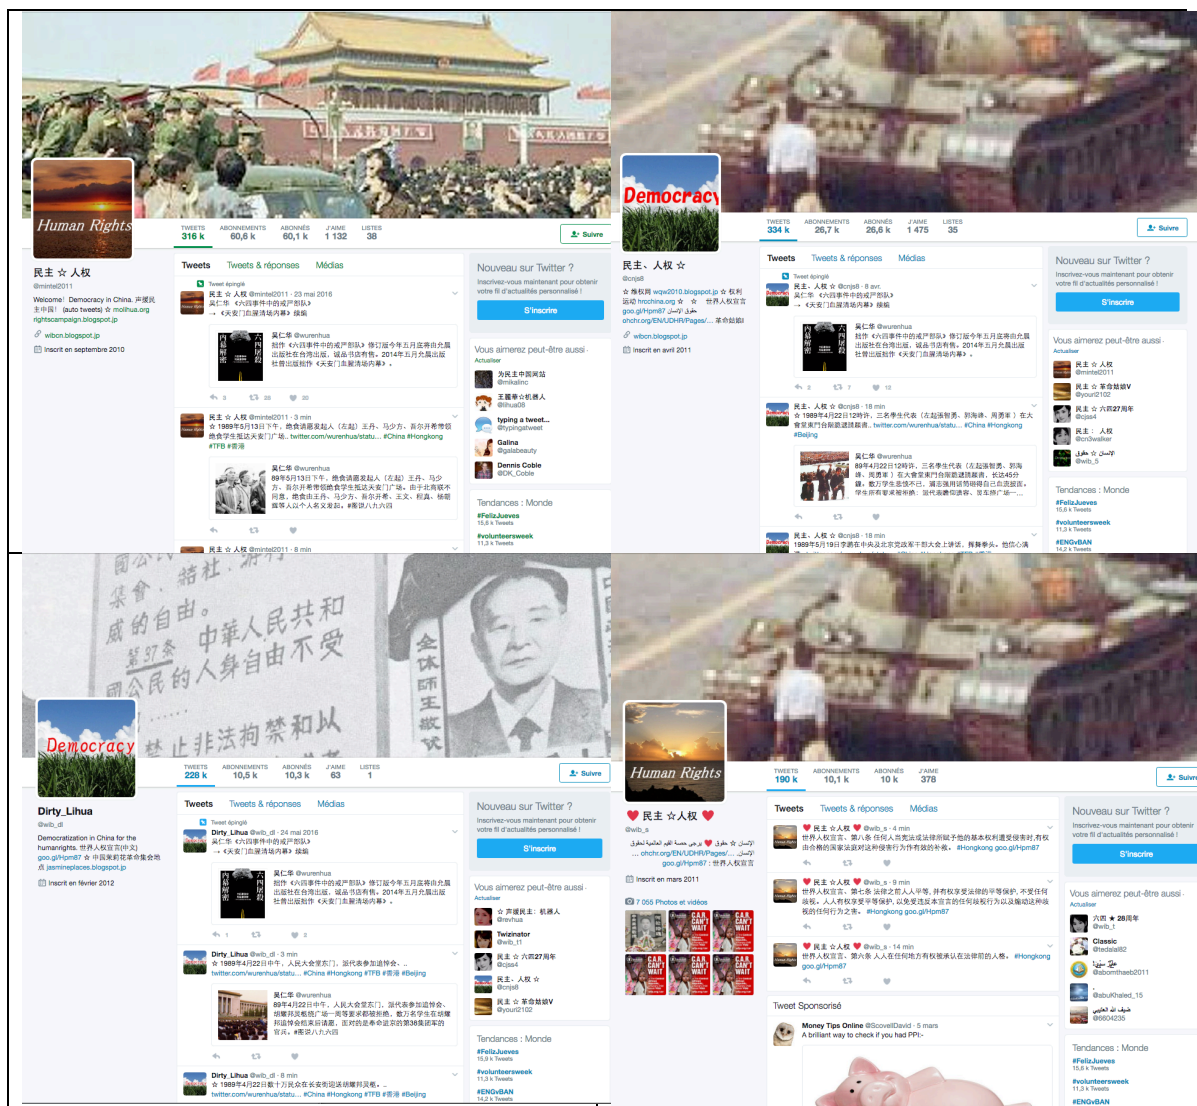
	Number of accounts in top 100 posters	Number of posts in dataset	Percentage of posts in dataset	Average BotOrNot Score
Anti-Chinese-state bots				
1989 group	22	117,578	9.98%	60
Pan-Asia group	22	44,678	3.79%	48
Independent anti-Chinese-state bots	5	7,969	0.68%	65
Both anti-Chinese-state and commercial content	1	1,090	0.09%	50
Other political bots				
Professional news bots	10	39,239	3.33%	48
“Fake news” bots	4	10,213	0.87%	71
Other non-political bots				
Commercial bots	8	34,860	2.96%	58
Job bots	6	8,592	0.73%	55
Other non-political bots	4	6,620	0.56%	39
Account suspended				
Account suspended	18	64,170	5.45%	
TOTAL	100	335,009	28.44%	

Accounts in this group often use variations on the same profile name “民主, 人权” (democracy, human rights). These accounts also use similar screen names (cnjs8, wib_dl, wib_s, cjss4, wib_z), similar profile pictures (often of generically attractive Asian women or photos with the words human rights or democracy), and similar or identical header pictures (images associated with human rights in China, such as the famous “tank man” in Tiananmen Square). Each of these 22 accounts posted, on average, 118 tweets per day in one of the monitored hashtags. These accounts all utilized twittbot.net, with 100% of their online activity conducted through this automation service.

Figure 1 shows the top four highest-posting accounts in this group and demonstrates their similarity. Three have almost identical screen names, two have identical profile pictures and two have identical header images. The profile pictures and

header images of all four accounts have a similar format. Three of the four accounts link to a blogspot.jp blog. While there is a variation in the number of friends and followers between these accounts, each of them has a very similar number of friends and followers, suggesting that they have gained followers through reciprocal following. Each of these accounts has posted at least twice in the previous 20 minutes.

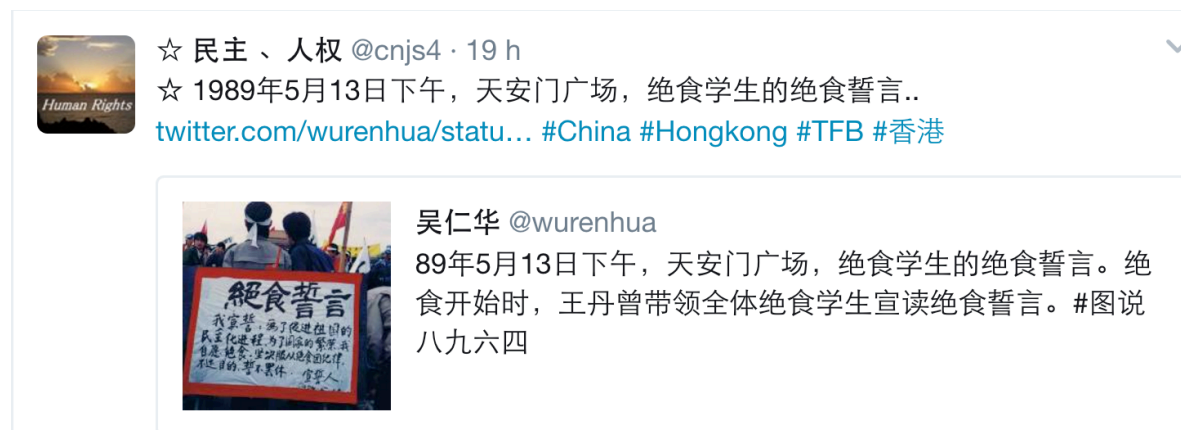
Figure 1. The top four highest-posting accounts in the 1989 bot group



The accounts in this group both post original content and retweet. All of the retweets were originally posted by 吴仁华 (@wurenhua), a leader in the 1989 movement who fled to America following the protests. Figure 2 shows two of these example posts.

Both of the original posts by wurenhua have a picture from the 1989 pro-Democracy movement. These bots retweet Wu Renhua's posts adding common hashtags to increase their dissemination.

Figure 2 Examples of forwarded posts from the 1989 bot group



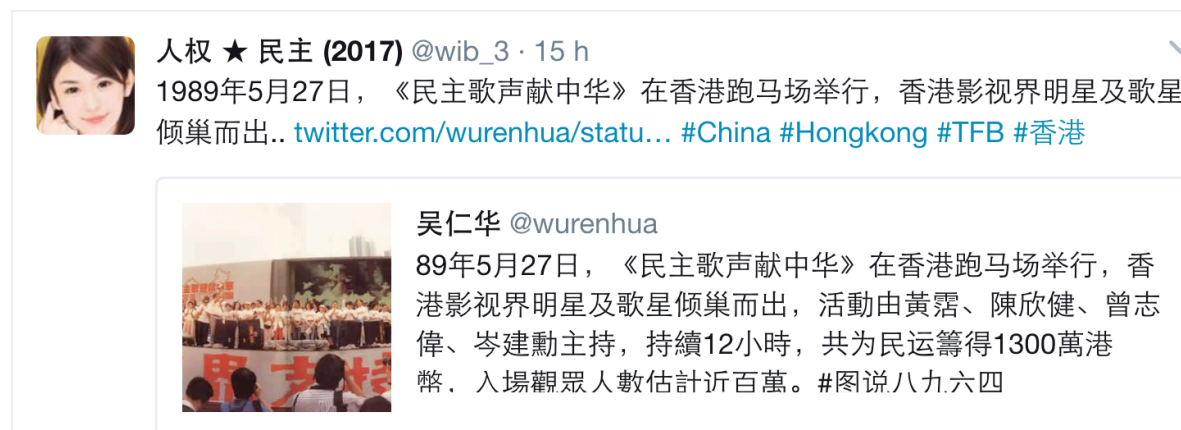
Translation:

☆ **Democracy, human rights @cnjs4 19 hours**

☆ On the afternoon of 13 May 1989 in Tiananmen Square, the students on hunger strike took an oath... <https://twitter.com/wurenhua/status/596489776821211136> ... #China #Hongkong #TFB #Hongkong

Wu Renhua @wurenhua

When the hunger strike began, Wang Dan led the hunger strike students to read the hunger strike oath. #Images of 4 June 1989



Translation:

Human rights ★ democracy (2017) @wib_3 15 hours

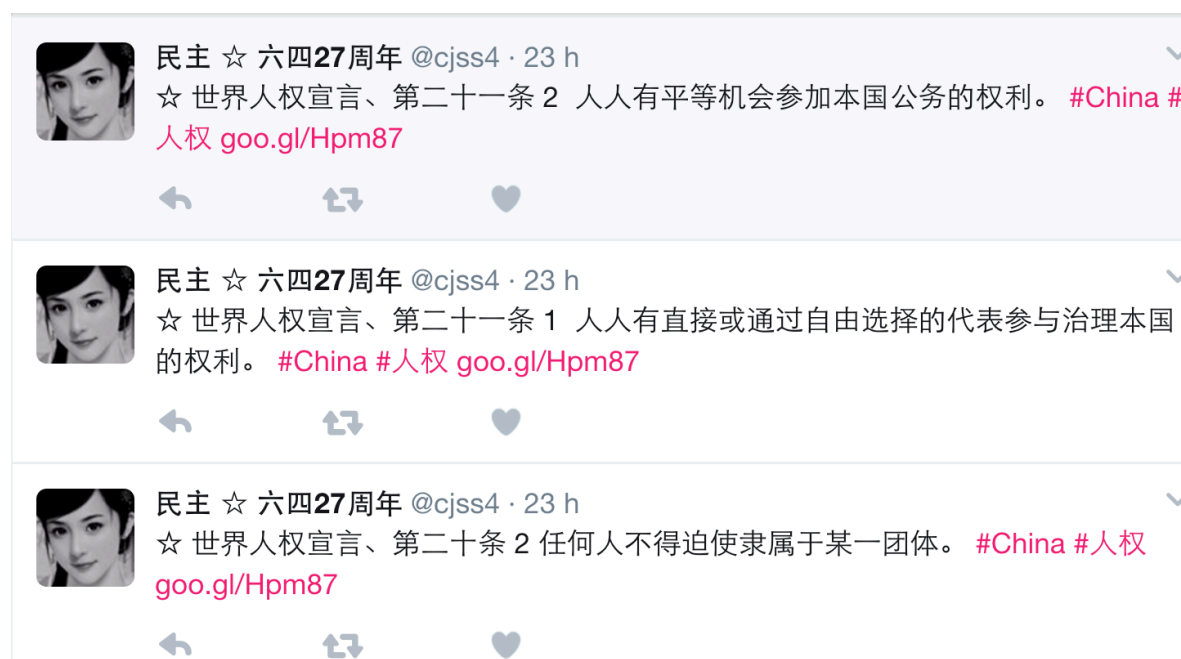
27 May 1989 "The Concert for Democracy in China" was held at the Hong Kong Racecourse, Hong Kong film stars and singers turned out in full force.... #China #Hongkong #TFB #Hong Kong

Wu Renhua @wurenhua

27 May 1989 "The Concert for Democracy in China" was held at the Hong Kong Racecourse, Hong Kong film stars and singers turned out in full force. The activities were presided over for 12 hours by Huang Zhan, Chen Xinjian, Eric Tsang and Cen Jianxun. A total of 13 million Hong Kong dollars was raised for the democracy movement and the number of viewers was estimated to be almost one million #Images of 4 June 1989

These bots also frequently post links to the Universal Declaration of Human Rights in Mandarin. All of these tweets were posted using the hashtags #China and #人权 (human rights); this means that, in particular, the hashtag #人权 is dominated by these bots. Eleven accounts in this group posted more than 1,000 times each using the hashtag 人权 during the data collection period, with the next highest poster posting 98 times. Almost 90% of the tweets that used the hashtag 人权 during the data collection period were posted by these 11 accounts. Figure 3 shows some example posts of this form and demonstrates how repetitive, formulaic and frequent these posts are.

Figure 3. Examples of original posts from the 1989 bot group



Translation:

Democracy ☆ 27th Anniversary of 4th June @cjss4 23 hours

The Universal Declaration of Human Rights, Article 21 2. Everyone has the right of equal access to public service in his country. #China #Human Rights [link to the Universal Declaration of Human Rights in Mandarin]

Democracy ☆ 27th Anniversary of 4th June @cjss4 23 hours

The Universal Declaration of Human Rights, Article 21 1. Everyone has the right to take part in the government of his country, directly or through freely chosen representatives. #China #Human Rights [link to the Universal Declaration of Human Rights in Mandarin]

Democracy ☆ 27th Anniversary of 4th June @cjss4 23 hours

The Universal Declaration of Human Rights, Article 20 2. No one may be compelled to belong to an association. #China #Human Rights [link to the Universal Declaration of Human Rights in Mandarin]

Given that the only previous reports of Chinese computational propaganda on Twitter have been of pro-state perspectives, the existence of this bot group is relatively surprising. This group is presumably aimed at the Chinese diaspora, students studying abroad, or those who jump the wall from the Chinese mainland to use Twitter. As a result information shared on Twitter with the hashtags commonly used by this bot group, such as #China and #人权 (human rights), appear to be dominated by this pro-democracy, anti-Chinese-state information. Indeed, this is not the only anti-state group posting in simplified Mandarin on Twitter.

The pan-Asia group

A second large group existed among the top 100 most frequently posting accounts in the dataset. This group disseminated information about the victims of the pan-Asia “Ponzi scheme.” Approximately 220,000 people lost the money they have invested in the Kunming Pan-Asia Nonferrous Metals Exchange when it collapsed in late 2015 (China Economic Weekly, 2015; VOA Chinese, 2015). There have been protests by those who lost money in this collapse and accusations that the local government was complicit in supporting the exchange.

This group appears to post less frequently than the 1989 group; the 22 accounts in this group that were among the top 100 posters in the dataset posted, on average, 43 times per day in one of the monitored hashtags. This is lower than the cut-off point of 50 tweets per day sometimes used to identify likely bot activity. The source of the tweets for accounts in this group are either Twitter for Android or Twitter for iPhone. Thus, although it is clear that this is a group of fake accounts, it is not clear that they are automated.

Many of the accounts in this group utilize similar screen names, such as GG8bjf0629Ehtvr, DkAvNtlRmLDHJYI and 5KMGRvJX9mSYaoQ. Several of the

accounts in this group present themselves as major Chinese news organizations or educational institutions in their display name, including 雲南日報 (Yunan Daily News), 中國新聞 (China News), 中國 • 瑞麗 News (China • Rili News), CCTV, 北京大学 (Peking University), 上海财经大学 (Shanghai University of Finance and Economics) and 吉林大学 (Jilin University)⁵. All of the accounts in this group listed their locations as being in the US.

Several of these accounts used the same information in their profile descriptions—despite being created at different times. For instance, the accounts named Jilin University (created in August 2016) and CCTV (created in February 2017) used an identical string of hashtags as their profile description: #China #Pan-Asia #Foreign Ministry #Travel #Nineteenth Party Congress #Xi Jinping #Pang Liyuan #Wang Qishan #Jiang Zemin #Meng Jianzhu #Beijing #Tiananmen Square #Peking University #Fudan University #Nanjing University #Wuhan University #Sun Yat-sen University #Xiamen University #Tsinghua University #Hong Kong university #United States #Trump #Harvard University #Cambridge University #University of Sydney.

Figure 4 shows an example of the posts of this group, which appear to predominantly retweet content published by other accounts in the group. Accounts in this group tweet with a wide number of hashtags. This group showed up frequently in the dataset for their use of hashtags such as #北京 (Beijing) and #习近平 (Xi Jinping). However, as Figure 4 shows, they also post frequently in hashtags that were not monitored as part of this data collection. Thus, more research would be necessary to uncover the true size of this group. However, what is clear is that automated and fake accounts that aim to disseminate information that attacks or is counter to the information

⁵ Surprisingly, despite publishing in simplified Mandarin (used in mainland China) many of the display names of accounts in this group utilised traditional characters: 雲南日報 instead of 云南日报 and 中國新聞 instead of 中国新闻. This suggests that this group might be linked with Taiwan, Hong Kong or Macau where traditional characters remain in use.

disseminated by the Chinese state are prominent in Chinese language information on Twitter. Indeed, these two groups are not the only fake accounts promoting anti-Chinese-state perspectives on Twitter.

Figure 4. Example of retweeted content in the pan-Asia group



Translation:

Shanghai University of Finance and Economics retweeted Stubborn Protest @juejiang01 3 May

The #Pan-Asia victims were forced to Lishan by the Kunming, Yunan government. This cannot be helped until suffering every possible torment they would go to Beijing to request national aid in hope of recovering justice and their hard-earned money. #Wang Qishan #Yao Ming #Meng Jianzhu #Xi Jinping #Central Commission for Discipline Inspection #Hainan Airlines #Guo Wengui @PDChina

Shangahi University of Finance and Economics retweeted Stubborn Protest @juejiang01 3 May

Kunming government documents set up Fanya to participate in fraud 43 billion. Pan-Asian Exchange. #Apollo #Wang He #Joan #Chang'an Street #Xinhua News Agency #Pan Asia

Other anti-Chinese-state bot activity on Twitter

This analysis also found evidence of other anti-Chinese-state bots (such as pro-Uighur and pro-Hong Kong independence bots) posting in simplified Chinese, Japanese and English. Restricting analysis to only hashtags associated with Tibet and Buddhism found no evidence of bots disseminating the pro-Chinese-state perspectives reported in the media in 2014. Instead, there was evidence of automation used to promote the messages of the Tibetan exile community and disseminate information about repression of ethnic Tibetans, predominantly in English. This analysis suggests that the Chinese state is not utilizing automation to influence discourse on Twitter. The implications of these findings for understanding Chinese international propaganda efforts are discussed in the conclusion section.

Computational propaganda on Weibo: little evidence of automation

In contrast to the high level of automaton in posts about China on Twitter, there was little evidence of automation in the Weibo dataset. Out of the 815,776 unique users in this dataset of 1,543,165 comments, only 145 users posted 100 or more comments across the examined posts. Based on an examination of their posting patterns, post content and post sources, these high-posting users did not appear to be using automation and there did not seem to be evidence that these were fake accounts.

However, the content of the posts of the highest-posting users indicates that there may be significant trolling within these comments. The majority of comments from the highest-posting user were attacks on other posters, which spanned multiple posts in the dataset. While the majority of users who posted comments on these stories appear to be genuine individuals posting their opinions and thoughts, this evidence of high posting by troll accounts would potentially drive the conversation away from productive discussions.

These findings are in line with previous research that found little evidence of automation in state-sponsored propaganda posts across a variety of platforms. Taken

together with the findings from Twitter, these results suggest that, perhaps surprisingly given the sophistication of the automated censorship functionality of the domestic Chinese Internet and the apparently wide use of automation by political interests in the US and Europe, automation does not appear to be being used as part of the Chinese state's propaganda strategy.

Conclusion

This article collected data to examine whether automation was present in hashtags associated with Chinese politics on Twitter and in comments on official news information on Weibo. These data indicate that the Chinese state is not using automation as part of either its domestic or international propaganda efforts. However, surprisingly, significant evidence of anti-Chinese state bot activity was found on Twitter, publishing predominantly in simplified Mandarin and presumably aimed at diasporic Chinese or those who 'jump the wall' to access foreign social media platforms.

While it may seem surprising to find that the Chinese state does not seem to be using automation, this can possibly be explained by several reasons. Firstly, Chinese international propaganda efforts have long been dominated by massive state-run companies such as CCTV, China Radio International and the China Daily. The focus on the Internet that intensified in 2016 has seen a rise in online media produced by traditional providers, such as the children's bedtime story explaining the One Belt, One Road policy posted to YouTube by the China Daily⁶ or the song about the 2016-2020 Five Year Plan posted to YouTube by China Global Network Television.⁷ Incorporating bots and automation into this international propaganda strategy would require new technological capabilities that are not the province of these traditional media providers. Thus, it may be the case that despite its technological sophistication and massive

⁶ https://www.youtube.com/watch?v=H6Adz_arAYE

⁷ <https://www.youtube.com/watch?v=LhLrHCKMqyM>

budgets, the Chinese state might be slow to incorporate bots into their propaganda strategy.

Secondly, bots and automation are a cheap and dirty solution to achieving particular ends; they allow single individuals or small groups to harness computational power to spread their messages more effectively. However, China is a strong state that can call on a massive supply of human resources. Thus, manually created and disseminated propaganda may be a smarter and more effective strategy. On the domestic Chinese Internet, research based on a leak from a local propaganda office found that, instead of the army of individuals paid 50-cents per post, Chinese online propaganda was mostly executed by state-employees acting as part of their regular jobs (King et al., 2017).

Similarly, a recent report on computational propaganda in Taiwan found that the examined incidents showed no evidence of automation or even state coordination; instead it was regular Chinese Internet users (albeit nationalistic ones), who seemed to be taking it upon themselves to promote reunification with China in the Taiwanese Internet sphere (Monaco, 2017). This suggests that rather than relying on bots, which would be subject to computational detection and whose functionalities are limited, the Chinese state can utilise its human resources both directly (by tasking state employees with posting positive information online) and indirectly (by cultivating and facilitating Chinese citizens influenced by domestic propaganda to promote Chinese-state interests both domestically and internationally).

This article uncovers the surprising fact that on Twitter (counter to media reports of Chinese state-associated bot activity) it is anti-state groups with few resources who are using automation to manipulate information about China and Chinese politics. One perspective on these results would be to conclude that Twitter and the use of automation on the platform is levelling the playing field for these less powerful voices to be heard.

However, when Chinese speaking users come to Twitter they are normally doing so because they want to find more diverse, less-biased information. They tend to see the platform as more akin to a public sphere, in contrast to China's more controlled online platforms. The fact that there is a great deal of automation, particularly within information in simplified Mandarin, suggests that Twitter is not acting as the kind of space for free information that these users hope to find.

It may be the case that influencing Twitter discourse about China in simplified Mandarin is not a priority for the Chinese state. Although Twitter use by mainland Chinese citizens is not as rare as its banned status might suggest, those who go out of their way to access foreign social media platforms are relatively likely to already hold anti-Chinese state perspectives. Targeting these Chinese Twitter users with pro-state propaganda would perhaps have little effect. However, these users would likely be susceptible to anti-Chinese state propaganda, supporting the existence of the bots uncovered in this article.

Twitter is also accessible to diasporic Chinese, including students studying abroad. However, information on the platform may have less effect on this population than might be hypothesized. Most Chinese students studying abroad continue to use domestic social media platforms such as Weibo, WeChat and QQ. It has also been reported that Chinese students who seen as holding anti-state views are denied visas or not selected for study abroad programs. Thus, the population of young Chinese who can access Twitter during their time abroad are already pre-selected as to be less susceptible to anti-Chinese state perspectives.

Another possible reason for the lack of Chinese state automaton on Twitter might be that these bots, in fact, have little effect. While this article and other similar studies, utilize hashtags to investigate the influence of bots on social media, prominence in hashtags does not necessarily translate into influence of discourse or opinions.

Information exposure on Twitter is primarily limited to information posted by accounts the user follows (and advertisements). As such, bot influence might be mostly limited to search results and trending topics. More research is necessary to investigate the influence of bots and bot-created content on public opinion.

Additionally, mostly in response to increased media and academic focus on online automation, social media platforms have committed to controlling bots. Thus, it is potentially the case that posts from bot accounts known to the platform would be prevented from appearing on user timelines and in search results. Previous research by the author on Weibo demonstrated that accounts and posts the user follows but that appear to be posting spam are hidden from user timelines (Bolsover, 2017). It would be reasonable to believe that Twitter also engages in a similar practice. Thus, more research is needed to uncover the true influence of bots on online discourse.

This research is also limited in several ways in several other ways. Firstly, the datasets are based on delineated time periods. It may be the case that automation is utilized surrounding particular events and the fast-moving nature of both the Internet and Chinese politics means that a lack of automation now does not necessarily mean a lack of automation in six months. Secondly, the conclusions of this article are based on posts in hashtags about Chinese politics on Twitter and comments on posts by official information providers on Weibo. Chinese state automation could possibly be found on these platform in other areas. On Twitter, Chinese state-associated automation could be being used to attack critics or foreign news organizations publishing in Chinese or to increase the dissemination of Chinese state-produced information. If these posts were not made during the timeframe examined using one of the hashtags examined, they would not be present in this dataset.

Thirdly, a conceptual limitation of this research is its focus on the use of bots and automation to achieve certain ends. As the case of Chinese domestic propaganda shows,

manual production and dissemination of online propaganda may be more effective than automated efforts. Given the extent of automation found in recent political events in the US and UK, continued research into bots on social media is important; however, the focus on automation should not blind researchers to the larger picture of online propaganda that includes cyborgs, hybrid accounts and manually produced propaganda.

Despite these limitations, this article provides the first academic insight into the use of automation to influence information about China and Chinese politics on international social media platforms. It also contributes to the limited knowledge about the use of bots on Chinese domestic social media. Perhaps surprisingly, given media reports of Chinese state-associated bots on Twitter, no evidence of Chinese state automation was found either domestically or internationally. This contributes to the literature on Chinese soft power and foreign diplomacy; despite indications that more attention would be paid to China's image on foreign social media, automation does not (yet) seem to be part of the country's international propaganda strategy.

Even more surprising was the finding of large amounts of anti-Chinese state automation in hashtags about China and Chinese politics on Twitter. While the true influence of bots on the beliefs and actions of social media site users is still unknown, almost 30% of the content in the examined hashtags was posted by bots. Very little is known about information on Twitter in the Chinese language or the way in which the platform might be being used to manipulate public opinion among Mandarin speakers.

The topic of automation, algorithms and online politics has only recently become a major area of investigation. This article is the first to address the question of the existence of computational propaganda about China on international social media and, thus, should not be the final answer to questions about this phenomenon. As research in this field progresses, it is important to remember that bots are not agentic nor are they isolated. They are created by individuals to fulfill specific functions. The concern about

bots and automation should not distract from the fact that these techniques are just a tool that is embedded in an underlying social structure. More focus should be paid to the political, social and economic systems that facilitate this kind of opinion manipulation and the conditions that mean their use is prevalent. More nuanced methods are also needed to detect online computational propaganda. Further efforts should move away from a solely computational and detection-based focus, to qualitative considerations of the content of automation-supported information to evaluate whether it is propaganda rather than whether it is simply computational. It is the first we are worried about not the second and this study has shown that the second is not always a proxy for the first.

References

- Bessi, A., & Ferrara, E. (2016). Social bots distort the 2016 US Presidential election online discussion. *First Monday*, 21(11).
- Bolsover, G. (In Press). It's political, 笨蛋: A random sample of microblog users and their speech acts in the US and China.
- Bolsover, G. (2013). News in China's New Information Environment: Dissemination Patterns, Opinion Leaders and News Commentary on Weibo (SSRN Scholarly Paper No. ID 2257794). Rochester, NY: Social Science Research Network.
- Bolsover, G. (2016). Harmonious communitarianism or a rational public sphere: a content analysis of the differences between comments on news stories on Weibo and Facebook. *Asian Journal of Communication*, 27(2), 1–19.
- Bolsover, G. (2017, January). Technology and political speech: commercialisation, authoritarianism and the supposed death of the Internet's democratic potential. University of Oxford, Oxford, UK.

- Bradshaw, S., & Howard, P. (2017). Troops, trolls and troublemakers: A global inventory of organized social media manipulation. COMPROP Working Paper 2017.12.
- China Economic Weekly. (2015, September 21). The Pan Asia crisis: A 40 billion “Ponzi scheme”? (泛亚危机：400 亿的“庞氏骗局”?). China Economic Weekly (中国经济周刊).
- Chow, E. K. (2017, June 29). China’s Propaganda Goes Viral. *The Diplomat*.
- Cobain, I., & Taylor, M. (2016, November 23). Far-right terrorist Thomas Mair jailed for life for Jo Cox murder. *The Guardian*.
- Collins, B., & Cox, J. (2017, October 27). This Twitter Bot Army Is Chasing Down a Chinese Dissident and Mar-a-Lago Member. *The Daily Beast*.
- Cook, S. (2015, December 10). Resisting Beijing’s Global Media Influence. *The Diplomat*.
- Creemers, R. (2015). Never the twain shall meet? Rethinking China’s public diplomacy policy. *Chinese Journal of Communication*, 8(3), 306–322.
- 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). ACM.
- Greitens, S. C. (2013). Authoritarianism Online: What can we learn from internet data in nondemocracies? *PS: Political Science & Politics*, 46(02), 262–270.
- Hassid, J. (2012). Safety Valve or Pressure Cooker? Blogs in Chinese Political Life. *Journal of Communication*, pp. 212–230.
- Hayden, C. (2012). *The rhetoric of soft power: Public diplomacy in global contexts*. Lexington Books.

- Helm, T. (2016, September 10). Brexit camp abandons £350m-a-week NHS funding pledge. *The Guardian*.
- Henochowicz, A. (2014, August 28). Murong Xuecun Targeted in Twitter Smear Campaign.
- Howard, P., Bolsover, G., Kollanyi, B., Bradshaw, S., & Neudert, L.-M. (2017). Junk News and Bots during the U.S. Election: What Were Michigan Voters Sharing Over Twitter? COMPROP Data Memo 2017.1.
- Howard, P., & Kollanyi, B. (2016). Bots, #StrongerIn, and #Brexit: computational propaganda during the UK-EU Referendum. COMPROP Research Note 2016.1.
- Howard, P., Kollanyi, B., & Woolley, S. (2016). Bots and Automation over Twitter during the US Election. COMPROP Data Memo 2016.4.
- Intelligence Community Assessment. (2017). Assessing Russian Activities and Intentions in Recent US Elections.
- Kaiman, J. (2014, July 22). Free Tibet exposes fake Twitter accounts by China propagandists. *The Guardian*.
- Kalathil, S. (2017). Beyond the Great Firewall: How China Became a Global Information Power. Center for International Media Assistance.
- King, G., Pan, J., & Roberts, M. (2017). How the Chinese Government Fabricates Social Media Posts for Strategic Distraction, not Engaged Argument. *American Political Science Review*, 111(3), 484–501.
- Li, B. (2013). Consciously persist in the Marxist news view. [Translation R. Creemers, China Copyright and Media]. Qiushi.
- Livingstone, S. (2016, January 7). Assessing China's Plan to Build Internet Power. China File.
- LoBianco, T. (2017, April 25). The 4 Russia investigations in Congress, explained. CNN.

- Monaco, N. (2017). Computational Propaganda in Taiwan: Where Digital Democracy Meets Automated Autocracy. COMPROP Working Paper 2017.2.
- Mustafaraj, E., & Metaxas, P. (2010). From obscurity to prominence in minutes: Political speech and real-time search. Proceedings of the WebSci10: Extending the Frontiers of Society On-Line, April 26-27th, 2010, Raleigh, NC: US.
- Ng, J. (2015, July 20). Tracking Censorship on WeChat's Public Accounts Platform. The Citizen Lab.
- Phillips, T. (2014, September 8). China's Communist Party takes online war to Twitter. The Telegraph.
- Qing, K., & Shiffman, J. (2015, November 2). Exposed: China's covert global radio network. Reuters.
- Ratkiewicz, J., Conover, M., Meiss, M., Gonçalves, B., Patil, S., Flammini, A., & Menczer, F. (2011). Truthy: Mapping the Spread of Astroturf in Microblog Streams. In Proceedings of the 20th International Conference Companion on World Wide Web (pp. 249–252). New York, NY, USA: ACM.
- Shambaugh, D. (2010, June 7). Opinion | China Flexes Its Soft Power. The New York Times.
- Shorey, S., & Howard, P. (2016). Automation, Algorithms, and Politics| Automation, Big Data and Politics: A Research Review. International Journal of Communication, 10(0), 24.
- Sullivan, J. (2012). A tale of two microblogs in China. Media, Culture & Society, 34(6), 773–783.
- VOA Chinese. (2015, September 22). The Pan-Asia Ponzi Scheme (泛亚“庞氏骗局” : 22 万投资人的家当有望讨回吗?). VOA Chinese (美国之音).

- Wang, Y. (2008). Public diplomacy and the rise of Chinese soft power. *The Annals of the American Academy of Political and Social Science*, 616(1), 257–273.
- Woolley, S. (2016). Automating power: Social bot interference in global politics. *First Monday*, 21(4).
- Xinhua. (2016, February 19). At the Party's News and Public Opinion Work Conference, Xi Jinping Stresses Persisting in the Correct Orientation, Innovating Methods and Means, and Raising the Dissemination Strength and Guidance Strength of News and Public Opinion. [Translation R. Creemers, China Copyright and Media]. *People's Daily*.
- Zhu, T., Phipps, D., Pridgen, A., Crandall, J., & Wallach, D. (2013). The Velocity of Censorship: High-Fidelity Detection of Microblog Post Deletions (arXiv e-print No. 1303.0597).