



Active, aggressive, but to little avail: characterizing bot activity during the 2020 Singaporean elections

Joshua Uyheng¹ · Lynnette Hui Xian Ng¹ · Kathleen M. Carley¹

Published online: 4 May 2021

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2021

Abstract

Digital disinformation presents a challenging problem for democracies worldwide, especially in times of crisis like the COVID-19 pandemic. In countries like Singapore, legislative efforts to quell fake news constitute relatively new and understudied contexts for understanding local information operations. This paper presents a social cybersecurity analysis of the 2020 Singaporean elections, which took place at the height of the pandemic and after the recent passage of an anti-fake news law. Harnessing a dataset of 240,000 tweets about the elections, we found that 26.99% of participating accounts were likely to be bots, responsible for a larger proportion of bot tweets than the election in 2015. Textual analysis further showed that the detected bots used simpler and more abusive second-person language, as well as hashtags related to COVID-19 and voter activity—pointing to aggressive tactics potentially fuelling online hostility and questioning the legitimacy of the polls. Finally, bots were associated with larger, less dense, and less echo chamber-like communities, suggesting efforts to participate in larger, mainstream conversations. However, despite their distinct narrative and network maneuvers, bots generally did not hold significant influence throughout the social network. Hence, although intersecting concerns of political conflict during a global pandemic may promptly raise the possibility of online interference, we quantify both the efforts and limits of bot-fueled disinformation in the 2020 Singaporean elections. We conclude with several implications for digital disinformation in times of crisis, in the Asia-Pacific and beyond.

Keywords Bots · COVID-19 Pandemic · Elections · Social cybersecurity

This work was supported in part by the Knight Foundation and the Office of Naval Research grants N000141812106 and N000141812108. Additional support was provided by the Center for Computational Analysis of Social and Organizational Systems (CASOS) and the Center for Informed Democracy and Social Cybersecurity (IDeaS). 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.

Extended author information available on the last page of the article

1 Introduction

Extensive studies tackle the large-scale efforts of inauthentic accounts like bots and trolls to sway public opinion on digital platforms (Ferrara et al. 2016; Shu et al. 2017). In this view, an important area of research concerns the development of computational tools to identify and characterize such information operations, constituting a growing multidisciplinary effort collectively known as social cybersecurity (Beskow and Carley 2019b; Carley et al. 2018a; Uyheng et al. 2019). Amidst an ongoing global pandemic, such research efforts have become crucial especially during high-profile events like national elections, where coordinated online activities bear potential to undermine democratic practice or exert foreign influence (Ferrara et al. 2020b; Uyheng and Carley 2020b). In this view, a shared concern in the field regards the extent to which information operations prove successful, how they achieve such success, and whether their dynamics can be understood in a systematic fashion (Shao et al. 2018; Varol et al. 2017).

Although literature in this area has grown significantly in recent years, much of the evidence base notably remains concentrated in Western contexts (Allcott and Gentzkow 2017; Ferrara et al. 2020a; King et al. 2020). Online disinformation, however, is a ubiquitous phenomenon across various geopolitical settings; regional differences challenge and refine normative assumptions about the nature and processes of online disinformation (Humprecht et al. 2020; Tapsell 2020b). For instance, studies in non-Western societies present unique challenges in a discipline where many tools are often biased toward English and the activities of Euro-American users, thereby requiring flexible, cross-cultural methodological pipelines (Cha et al. 2020; Uyheng and Carley 2020a), and accounting for broader cultural practices and the local political setting (Tapsell 2020a; Udupa 2018).

Lying at the critical intersection of these concerns, Singapore held its general election to select members of its 14th Parliament in July of 2020. Although the incumbent People's Action Party had held overwhelming majority power for decades, this year's campaigns featured rising tides of support for the progressive and youth-oriented Worker's Party (Zaini 2020). Given pandemic-related disruptions and restrictions, national polls likewise took place at a volatile time, requiring careful health and safety mandates and shifting the norms of electoral campaigning (James 2020). Meanwhile, in the year prior, Singapore had legislated the Protection from Online Falsehoods and Manipulation Act (POFMA), which imposed large penalties for engagement in fake news propagation (Tan 2020). Criticized for its potential weaponization against democratic dissent and freedom of speech (Han 2019), the recent enactment of POFMA presented an additional layer of complexity to public discourse within an overarching setting of political contests during a global health crisis.

Amid pandemic-fueled economic uncertainty, did generational shifts in political sentiment create an environment ripe for online disinformation? Or did stringent legislation dampen widespread interference in online discourse? Adopting the lens of social cybersecurity, we empirically tackle these questions through a descriptive analysis of Twitter conversations surrounding the 2020 Singapore elections.

Although we do not advance strict causal claims in this work, we show that bot activity has increased significantly since the previous elections in 2015, and quantify their aggressive activities from both a narrative and network perspective. From a methodological standpoint, we also show how generalizable tools harnessing machine learning and network science can be used within a uniquely diverse linguistic, political, and cultural setting. We also present practical insights for worldwide elections more broadly. In sum, we therefore ask the following questions:

1. How prevalent were Twitter bots in the 2020 Singaporean elections?
2. How can we characterize the activities of Twitter bots in the 2020 Singaporean elections?
3. How influential were the Twitter bots in the 2020 Singapore elections?

2 Related work

2.1 Digital democracy and online disinformation

It has become a matter of broad consensus that the advent of digital platforms like social media has transformed public discourse. Large-scale and relatively unfettered connectivity between individuals online has enabled and accelerated the formation of digital communities and innovative pathways for political participation (Gil de Zúñiga et al. 2010). Alongside constructive views of ‘embracing’ the democratic potentials of cyberspace (Gastil and Richards 2017; Noveck 2017), many scholars also point to the interconnected issues of fake news, political polarization, and echo chambers which arise as a combined result of the social and technological dynamics of these platforms (Allen et al. 2020; Geschke et al. 2019; Tucker et al. 2018).

Taking center stage amidst these issues is the problem of information operations, which define concerted campaigns to take advantage of these sociotechnical factors in pursuit of disruptive objectives (Beskow and Carley 2019b; Shallcross 2017). One important ingredient in such campaigns is the social bot, which refers to agents on digital platforms relying on some level of automation to achieve a strategic goal (Caldarelli et al. 2020; Ferrara et al. 2016). Social bots have been documented in a wide range of major high-profile issues with national and global consequences, especially in the 2016 and 2020 US elections (Allcott and Gentzkow 2017; Ferrara et al. 2020a), Brexit in the UK (Bastos and Mercea 2019), elections around the world, (Keller and Klinger 2019; King et al. 2020; Uyheng and Carley 2020a), and with connection to state-sponsored activities (Badawy et al. 2018; Uyheng et al. 2019; Woolley 2016). They therefore represent a significant concern in securing the free and open nature of cyberspace in societies worldwide.

2.2 Social cybersecurity during crisis

In response to this wave of digital threats, the burgeoning field of social cybersecurity opens a multidisciplinary and multi-methodological space for tackling these

issues (Carley et al. 2018a; Carley 2020). By adopting a sociotechnical view, a social cybersecurity lens mobilizes models and insights around digital disinformation in computer science (Beskow and Carley 2018; Shu et al. 2017) to examine its dynamics and impacts within the social realm (Starbird et al. 2019; Uyheng and Carley 2020a). Operationally, this translates into a consideration of *narrative maneuvers*, which may involve either increasing positive or negative storylines around a person, event, or organization; as well as *network maneuvers*, by which social bots may shift the flow of information through the formation, connection, and dismantling of group structures (Beskow and Carley 2019a; Şen et al. 2016).

Integrated perspectives such as those of social cybersecurity are particularly salient in the context of global crises like the COVID-19 pandemic. Social scientists have laid wide-ranging foundations for the use of the social and behavioral sciences to aid in pandemic response (Van Bavel et al. 2020). Meanwhile, social media scholars and computer scientists have also characterized the so-called ‘infodemic’ associated with the massive spread of information related to COVID-19 (Cinelli et al. 2020). Concerns over digital disinformation lie at the intersection of both issues, whereby information campaigns take advantage of pandemic-related chaos to sow discord and spread falsehoods on digital platforms (Ferrara et al. 2020b). Infodemics have moreover taken on a particularly hateful character during the global crisis, with targeted aggression online directed especially toward political leaders and marginalized social groups in response to broad societal strain (Uyheng and Carley 2020b, 2021).

2.3 Contexts of disinformation in the Asia-Pacific

It is against this backdrop that we analyze the 2020 Singaporean elections, which represented a major national political process taking place at the height of the pandemic. One important consideration here is its broader geopolitical context in the Asia-Pacific, in contrast to much of the foregoing literature concentrated primarily in the West (Allcott and Gentzkow 2017; Ferrara et al. 2020a; King et al. 2020). We situate our work in relation to wider calls to improve comparative and cross-cultural work in the study of digital disinformation, in alignment with its similarly global reach and international consequences (Bradshaw and Howard 2018; Humprecht et al. 2020).

Related work in the region demonstrates that despite the predominant emphasis on Russian disinformation in the West, similar issues pervade societies in the Asia-Pacific (Cha et al. 2020; Dwyer 2019). On a fundamental level, issues of ‘post-truth’ politics and fake news have also become salient in the continent (Yee 2017). Political participation through social media activity has likewise become a norm during elections, constituting a dual face of electoral campaigning alongside more formal practices (Tapsell 2020b).

But beyond the incorporation of digital platforms into local electoral politics, prior work has also observed that at most between 9.89% and 21.08% of participating users were predicted to be bots leading up to national elections in the Philippines, Indonesia, and Taiwan (Uyheng and Carley 2020a). Studies of online blogs

likewise show significant levels of toxicity over time and across various national settings in the region (Marcoux et al. 2020). Collectively, this suggests that national elections in the region have historically been vulnerable to information operations, which in Singapore, may be complicated by the unique contexts of the COVID-19 pandemic, in conjunction with the stringent POFMA legislation (Han 2019; Tan 2020).

2.4 Contributions of this work

This work empirically quantifies the presence of bots and analyzes their activity in the 2020 Singaporean elections. The 2020 Singaporean elections represent an important geopolitical event at the intersection of the COVID-19 pandemic, the POFMA legislation, and the context of the Asia-Pacific. Our findings therefore add to the growing global literature on digital disinformation, especially in times of crises, with potential insights into the consequences of ‘fake news’ laws. Finally, we demonstrate a flexible and generalizable methodology for operationalizing a social cybersecurity lens. As we outline in the succeeding sections, our use of machine learning and network science tools and methods enables systematic analysis of narrative and network maneuvers in a manner that is applicable to a variety of social issues beyond those we examine here.

3 Data and methods

3.1 Data collection

To sample the online conversation about the 2020 Singaporean elections, we employed Twitter’s REST API. Data collection was performed on a daily basis using general hashtags related to election discourse (e.g., #GE2020, #sgelections2020) and more specific terms related to prominent parties (e.g., @PAPSingapore, @wpsg) and parliamentary candidates (e.g., @jamuslim). Search terms were updated and validated based on manual searching of ongoing Twitter conversations; new hashtags were added to the search parameters as they emerged. Data collection began June 18—a week before the previous parliament was dissolved for the elections—and concluded July 17—a week after election day. A total of 240K tweets were collected featuring 42K unique users.

To provide a benchmark for local bot activity in Singapore, we additionally sampled the online conversation on the 2015 Singaporean elections. Due to standard restrictions in the Twitter API, we were unable to perform similarly comprehensive data collection for the prior election. Instead, we hydrated tweets from two previously collected datasets, obtained independently by two separate analysis efforts with the election hashtag #GE2015. Similar to the present work, the time frame of both datasets spanned the day after Parliament was dissolved to the day after the general elections (Chong 2015; Daryani 2015). These datasets are named 2015-A and

Table 1 Summary of datasets used for analysis of bot activity during 2020 Singaporean elections relative to 2015 Singaporean elections as baseline

Dataset	No. of Tweets	No. of Users	Source
2020	240K	42K	Authors
2015-A	3K	1.5K	Daryani (2015)
2015-B	1K	~300	Chong (2015)
2015-Dates	4K	2.4K	Authors

2015-B. We also performed a Twitter search on the same hashtag (#GE2015) constrained to the month of September 2015, producing a dataset named 2015-Dates.

Although these latter datasets are relatively small, they provide the most readily available snapshot of the previous election to serve as a provisional baseline for present bot activity. Table 1 summarizes key features of the datasets used in this work.

3.2 Bot detection

We used the BotHunter algorithm to identify inauthentic accounts in our dataset. BotHunter is a machine learning algorithm based on a random forest model trained on a large dataset of known bots (Beskow and Carley 2018). Employing a tiered approach, BotHunter utilizes account and network features to generate probabilistic predictions of whether an account is bot-like or not. BotHunter also features comparable predictive performance to existing bot detection algorithms in the literature, as demonstrated through rigorous comparative analysis on multiple benchmark datasets (Beskow 2020).

3.3 Messaging analysis

To characterize bot messages, one tool we used was Netmapper¹. A rich tradition in psycholinguistics and social psychology associates the use of particular words and expressions with behavioral, cognitive, and emotional states (Pennebaker et al. 2003; Tausczik and Pennebaker 2010), as well as deceptive and persuasive communication (Addaood et al. 2019; Pérez-Rosas et al. 2018). Using the Netmapper software, we count the frequency of key lexical categories including abusive terms, absolutist terms, exclusive terms, and positive and negative terms (Carley et al. 2018b). Netmapper is particularly useful in the Singaporean setting given its multilingual functionality, covering over 40 languages. Harnessing these measures, we specifically sought to distinguish the language employed by bot and human accounts.

In addition to this psycholinguistic analysis, we performed hashtag analysis to further characterize messaging on Twitter. Acknowledging that hashtag usage can vary dramatically by raw scale, we obtained the ordinal ranking of hashtag usage for predicted bot and human accounts. We were particularly interested in hashtags which ranked more highly for bots than for humans, indicating disproportionate

¹ <http://netanomics.com/netmapper/>

bot-driven focus on a particular message beyond the rest of the baseline conversation. In the context of an election during the pandemic, we also sought to assess the prevalence of hashtags related to COVID-19 as well as to voting in general. Hence, we examined the relative ranks of all hashtags containing the (case-insensitive) strings ‘vote’ or ‘covid’.

3.4 Social network analysis

Finally, we used ORA for social network analysis (Carley et al. 2018b). ORA² is an integrated tool for the analysis of large-scale, complex networks. We represented our Twitter corpus as a complex graph structure which contained multiple types of nodes—agents, tweets, hashtags—featuring multiple types of edges, including: agent by agent communication through retweets, replies and mentions; agent by tweet relationships based on who send what; agent by hashtag usage networks; and tweet by hashtag connections based on the hashtags contained per tweet.

These network structures enable a wide variety of pertinent analysis for social media conversations, including the measurement of user influence and the automatic detection of emergent community structure. For the former, numerous measures of centrality abound for assessing different notions of user importance within a network structure (Carley et al. 2018b). For the latter, we use the Leiden clustering algorithm, a known improvement over the Louvain clustering algorithm, which boasts mathematical guarantees for non-degenerate groups and faster run-time (Traag et al. 2019). Table 2 summarizes all the relevant network measurements employed on both the cluster and agent level for this work.

4 Results

Our findings reveal significant bot activity on Twitter surrounding the 2020 Singaporean elections. Using our interoperable pipeline of social cybersecurity tools, we further present a nuanced picture of distinct bot behaviors as well as evidence that their influence over the online conversation was relatively low.

4.1 Bot prevalence and interactions

Figure 1 depicts the network of users in our dataset based on their combined retweets, replies, and mentions. The proportion of red nodes (which represent bots) suggests a significant number of bots participating in the Twitter conversation. They also appear relatively ubiquitous throughout the social network. Interestingly, a manual inspection shows that a significant number of bots are personal

² <http://netanomics.com/ora-pro/>

Table 2 Summary of cluster-level and agent-level network features measured with ORA to quantify structural properties of communities and individual network influence

Level	Measurement	Description
Cluster	Size	Number of agents in the cluster
	Density	Proportion of actual over possible interactions in the cluster
	Cheeger Value	Extent of bottleneck behavior in the cluster (Mohar 1989)
	E/I Index	Normalized difference between external and internal cluster communication (Krackhardt and Stern 1988)
	Echo Chamberness	Weighted association between density of agent by agent communication network and agent by agent shared hashtag network (Carley et al. 2008)
	In-Degree	Sum of incoming links indicating extent to which agent retweets, mentions, or replies to others
	K Core	Size of maximal group of actors connected to k others in the group (Dorogovtsev et al. 2006)
	No. of Followers	Number of accounts following the agent
	Out-Degree	Sum of outgoing links indicating extent to which agent is retweeted, mentioned, or replied to
	Page Rank	Extent to which account interacts with other influential accounts (Brin and Page 1998)
	Produced Mentions	Number of mentions produced by the agent
	Received Mentions	Number of mentions received by the agent
	Total Degree	Total number of interactions received and produced by the agent




Fig. 1 ORA visualization of Twitter conversation surrounding 2020 Singaporean elections. Accounts are represented as nodes connected by edges weighted by the sum of all retweets, replies, and mentions. Node colors are red if BotHunter probability is greater than 0.8, and blue otherwise

informational bots, links with services like IFTTT (If This Then That) to retweet posts from either Twitter or other social media platforms.

Figure 2 quantifies these observations, showing that the distribution of BotHunter probabilities is skewed to the left, with a mean value of 0.62. At a 0.8 probability threshold for bot-likeness, 26.99% of unique users in our dataset may be classified as bots. Figure 2 further suggests that bots did not only constitute a large proportion of the users participating in the conversation; more bot-like accounts also produced more tweets on average. We observe a clear upward trend, with the most bot-like accounts producing about twice as many tweets as the most human-like accounts.

However, bots generally performed fewer interactions than humans, despite producing more tweets on average. This suggests bots produced many original tweets not directed at others. Furthermore, both bots and humans talk to humans more, as 56.09% of all interactions took place between humans, while 6.40% of all interactions came from bots directed at humans. But 32.66% of all communication by humans was also directed at bots. Whether knowingly or not, humans thus frequently communicated with bots.

In comparison with the 2015 elections, Figure 3 presents interesting—if mixed—findings. First, we notice that the proportion of tweets by bots is much

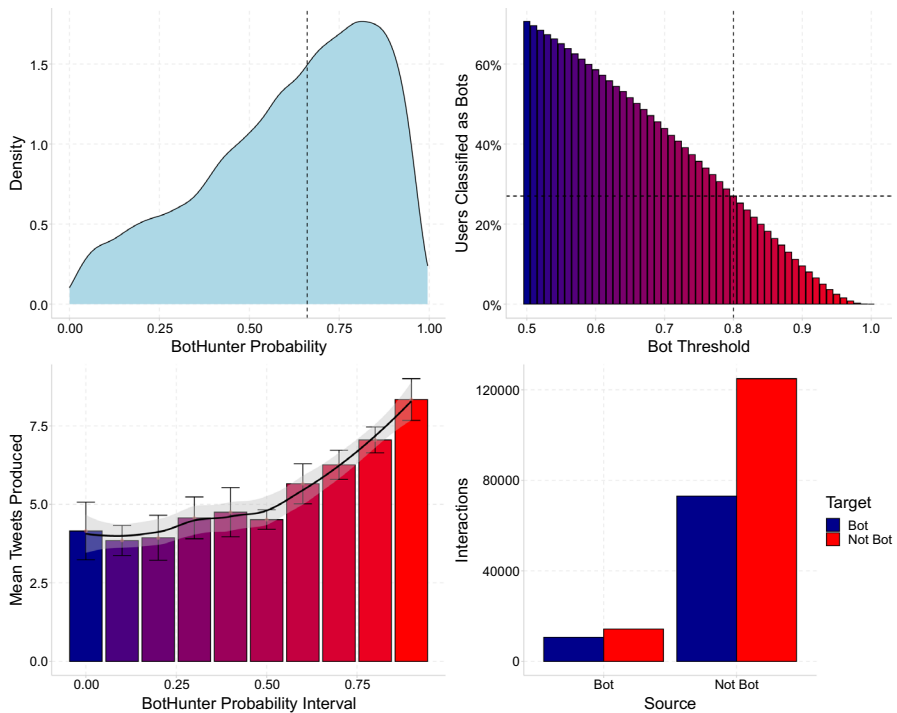


Fig. 2 Bot prevalence and behavior. Top-Left: Density plot of unique users' bot probabilities using BotHunter. Vertical line indicates mean value of bot probabilities at 0.62. Top-Right: Bar plot indicating percentage of bots at different BotHunter probability thresholds. At a 0.8 threshold (vertical line), 26.99% of unique users are classified as bots (horizontal line). Bottom-Left: Mean number of tweets produced by users at different intervals of BotHunter probabilities in 0.1 increments. Error bars represent 95% confidence intervals with fitted loess trend. Bottom-Right: Number of interactions between bots and humans based on a 0.8 probability threshold

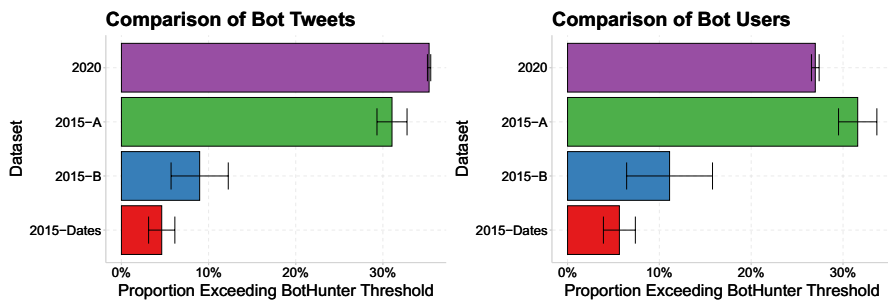


Fig. 3 Comparison of 2020 dataset with various datasets on 2015 Singaporean elections. Left: Comparison of proportions of tweets by bots. Right: Comparison of proportions of bot users




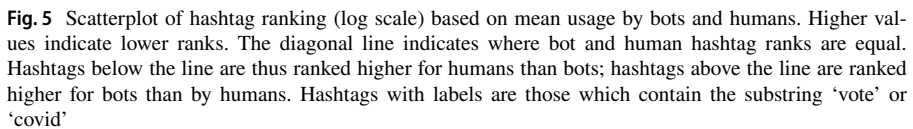
Fig. 4 Psycholinguistic cues which distinguish bots and humans. Points represent coefficient estimates in a multiple regression model predicting bot probability based on lexical features. Error bars represent 95% confidence intervals. Intersection of confidence intervals with the origin (broken line) indicates non-significant effects

higher in 2020 than in 2015, regardless of the dataset collection method. However, when comparing the proportion of bot users, one dataset (2015-A) achieves a significantly higher proportion of bot users. Although the other two datasets from 2015 do not offer similar measurements, the most conservative picture presented by our analysis is that bots may have reduced in relative proportion as actors between 2015 and 2020, but their domination of the electoral conversation has nonetheless increased in share. Notwithstanding the relatively small datasets used to obtain these snapshots, as well as the likelihood that Twitter may have already suspended bots from these previous datasets prior to rehydration, these findings point to consistent evidence of notable bot activity in the present elections above and beyond the only available evidence from prior elections.

4.2 Narrative maneuvers: bot aggression and electoral distrust

Besides interaction patterns, the language used by bots also featured notable differences from humans. Figure 4 shows the coefficients of a regression model relating BotHunter probability scores to users' psycholinguistic cues. All variance inflation factors ranged from 1.01 to 1.27, indicating no multicollinearity problems.

Most notably, we saw that bots used much simpler language than humans as denoted by the Reading Difficulty score. Human accounts were also significantly more likely to refer explicitly to identity terms related to gender or politics, as well as sentiment terms in general, whether positive or negative. In contrast, we found that bots were more likely to use abusive terms, absolutist terms, exclusive terms, and second-person pronouns. This suggests that, on average, bots were consistently engaged in more insulting behavior directed toward the people they interacted with.



In analyzing more specific patterns of hashtag usage, we point to three key observations. Figure 5 plots all hashtags used by bots and humans, ranked according to their total usage by both account types. First, we notice that most hashtags cluster around the diagonal line, which indicates where hashtags have the same ranking for bots and for humans. Many high-ranking hashtags for humans were also high-ranking for bots. Conversely, many low-ranking hashtags for humans were also low-ranking for bots. From this standpoint, we therefore note that many features of the online conversation were prioritized by both bots and humans. This intuition is borne out by a Spearman’s correlation test, which results in statistical significance ($\rho = 0.5926$, $p < .001$).

Finally, we consider the hashtags *above* the diagonal line, where we observe several hashtags which had *high rank for bots, but low rank for humans*. Most notably,




Fig. 6 Structural features distinguishing Leiden clusters with high or low bot activity. Points show coefficient estimates in a multiple regression model of average bot probability based on cluster structure. Error bars represent 95% confidence intervals

we see in relation to ‘vote’ the notion of ‘#votersuppression’. Conversations containing this hashtag cast suspicion on the integrity of elections during the pandemic, particularly by pointing to the dangers of the disease, as well as linking to related issues in the United States. Meanwhile, additional partisan hashtags (e.g., ‘#VoteThemOut’)—including ‘#covidiot’, which was directed against the PAP—fueled additional hostility versus the incumbent party. Collectively, these patterns suggest that bots—in contrast to humans more associated with mainstream messaging—were more likely to share messages related to the (un)trustworthiness of the elections under the shadow of the pandemic, or to fuel more incendiary political conversations. From a social cybersecurity standpoint, these behaviors suggest negative narrative maneuvers aimed at potentially reducing trust in the elections, as well as stoking enmity with the political faction currently in power.

4.3 Network maneuvers: bot cluster density and community interference

Bots operate not just through the messages they send, but also through manipulating social network structure based on artificial patterns of interaction with other users. Based on the results of Leiden clustering, we therefore asked: What structural features distinguish clusters with higher bot activity?

Regression analysis visualized in Figure 6 indicated that, on average, clusters featuring high levels of bot involvement simply tended to be larger, less dense, and less echo chamber-like. This suggests that bots tended to operate where the conversation was most active and communicated with people talking about a variety of topics. From the lens of social cybersecurity, these findings potentially point to positive network maneuvers, whereby inauthentic actors attempt to build groups or push information to larger—and more diverse—communities. The Cheeger value and E/I index of clusters did not predict levels of bot activity beyond the other factors already mentioned; hence, bots did not particularly exhibit hierarchical structure in clustered interactions, or preference for more or less isolated clusters. All variance

Table 3 Most influential accounts in 2020 Singaporean elections. A '+' indicates verified accounts; a '*' indicates news accounts

Rank	Super Spreaders	Super Friends	Other Influencers
1	wpsg ⁺	eisen	Reuters ^{+,*}
2	ChannelNewsAsia ^{+,*}	historyyogi	Cristiano ⁺
3	plspreeti ⁺	tanhuiyi	leehsienlong ⁺
4	jamuslim ⁺	sselection	wpsg ⁺
5	tzebern_	kixes ⁺	narendramodi ⁺
6	historyyogi	guanyinmiao	historyyogi
7	eisen	mdzulkar9	jamuslim ⁺
8	RaeesaKhanwpsg ⁺	plspreeti ⁺	nytimes ^{+,*}
9	mediumshawn	mediumshawn	sselection
10	MothershipSG ⁺	BenChiaCars	fat__thin

inflation factors in this regression model ranged from 1.20 to 3.55, indicating minimal issues with multicollinearity.

4.4 Bot failure to amass network influence

Finally, we consider the level of influence bots had relative to other users in the dataset. ORA summarizes super spreaders, super friends, and other influencers in the online conversation (Carley et al. 2018b). *Super spreaders* generate highly shared content, measured by average ranking on out-degree centrality (many share their content), page rank centrality (they interact with other influential accounts), and large k-core membership (belong to large cluster). *Super friends* engage in extensive two-way communication, again identified by highest average ranking on total degree centrality (total interactions) and large k-core membership. *Other influencers* are influential in other ways, by having high numbers of followers, or high levels of mentioning and being mentioned.

Table 3 provides the top 10 accounts for each of the three categories. Generally, it seems that the most influential accounts are verified accounts and news accounts. The opposition Workers Party, in particular, dominates the list by having their party account and individual candidates feature as super spreaders and other influencers. Incumbent Prime Minister Lee Hsien Long only appears among the top 10 other influencers. Bot accounts did not occupy dominant positions in these influencer lists either.

Figure 7 utilizes boxplots to visualize the distribution of influence metrics used by ORA relative to bot probabilities in 10% intervals. In all measures but one, bots were not more influential. Most bots did not belong to larger k cores, did not have more followers, did not have more viral content (out-degree), did not interact with more influential accounts (page rank), and neither received nor produced the most mentions and total interactions. However, bots tended to have higher in-degree, indicating that they produced more retweets, replies, and mentions in attempts to gain influence, but not necessarily securing it in a meaningful way throughout the larger conversation.

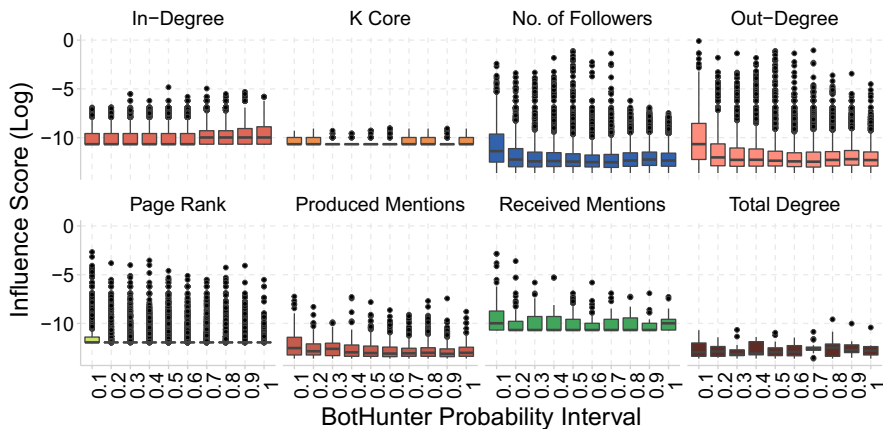


Fig. 7 Boxplots of influence scores (log) versus BotHunter probability intervals

5 Conclusions and future work

We reflect on four implications of our quantitative portrait of Twitter disinformation during the 2020 Singaporean elections. First, we highlight evidence of disruptive bot messaging related to COVID-19 and voter suppression. Notwithstanding genuine concerns the pandemic poses for equitable elections, researchers may examine how future information operations use these concerns to undermine democratic practice. Although evidence is scant that these messaging tactics had practical impacts during the election proper—especially given bots’ relatively low influence—these findings demonstrate how information operations harness wider societal conditions to potentially sow discord and undermine democratic integrity through the use of both narrative and network maneuvers (Carley 2020; Cinelli et al. 2020; Ferrara et al. 2020b). Salient bot-fueled narratives, in particular, appeared tailored toward distrust against political leaders and ‘othering’ them as violators of COVID-19 behavioral guidelines. Such identity-based dynamics, social scientists have pointed out, are factors at the core of social cohesion and collective behavioral efficacy against the pandemic (Van Bavel et al. 2020). Hence, here we find that not only does the global crisis occasion specific forms of online political attacks, but the aggressive, potentially bot-amplified nature of electoral contests may also have adverse impacts on the public’s response to the pandemic (Uyheng and Carley 2020b, 2021).

Second, we consider the effects of POFMA (Han 2019; Tan 2020). Although we did not explicitly find fake news sharing, we detected bot signals comparable with other Asia-Pacific countries without strict legislation (Uyheng and Carley 2020a). While we reiterate that causal claims are outside the scope of this work, the high prevalence of bots in online electoral discourse points to ways that digital disinformation—and online communities more broadly—may adapt to existing legislation to attempt other strategies of disruption, even if relatively unsuccessful. Despite the absorption of POFMA into the public consciousness (Ng and Yuan 2020), this raises questions about disinformation’s flexibility—exceeding falsehoods to include

hostility and discord—and effective ways of curbing it without curtailing free speech.

Third, we observe that the opposition held more online influence than the incumbent, but the latter still won parliamentary majority. This is consistent across recent elections in the Philippines, Indonesia, and Taiwan (Uyheng and Carley 2020a), suggesting that regional links between online popularity and electoral success are far from straightforward. In view of the strong social media component of electoral campaigns both in the Asia-Pacific and worldwide (Gil de Zúñiga et al. 2010; Tapsell 2020b), these findings suggest new points of inquiry for understanding the potentials (and limits) of online political participation in relation to eventual electoral outcomes. This insight may also be interrogated within the broader sphere of global studies of digital disinformation, which may rely on assumptions regarding the correspondence between online and offline sentiments which do not play out neatly here.

Fourth and finally, we affirm the importance of designing interoperable pipelines for social cybersecurity (Uyheng et al. 2019). Although it is not our goal here to innovate new tools in their own right, we show how the problem-oriented integration of existing tools can quantify unique narrative and network features of disinformation actors (Beskow and Carley 2019b; Carley et al. 2018a), and provide (negative) evidence of their influence over the broader conversation.

Several limitations nuance our conclusions from this work. Sampling Twitter data remains limited by API generalizability issues, suggesting nuance in extrapolating findings to wider contexts (Morstatter et al. 2013). We also reiterate our primarily empirical, rather than methodological, goals in this research. To improve these computations, algorithmic developments may consider local patterns of language and social media use (Cha et al. 2020). For instance, bot activity driven by malicious state actors bears different social significance from potentially more benign automation, such as those by services like IFTTT we observed here. Multi-platform studies would also aid more holistic inquiry into online electoral discourse given that Twitter may not play the same role everywhere as in the West (Tapsell 2020a).

References

- Addawood A, Badawy A, Lerman K, Ferrara E (2019) Linguistic cues to deception: Identifying political trolls on social media. In: *Proceedings of the International AAAI Conference on Web and Social Media*, vol 13, pp 15–25
- Allcott H, Gentzkow M (2017) Social media and fake news in the 2016 election. *J Econ Perspect* 31(2):211–36
- Allen J, Howland B, Mobius M, Rothschild D, Watts DJ (2020) Evaluating the fake news problem at the scale of the information ecosystem. *Sci Adv* 6(14):eaay3539
- Badawy A, Ferrara E, Lerman K (2018) Analyzing the digital traces of political manipulation: The 2016 Russian interference Twitter campaign. In: *2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, IEEE, pp 258–265
- Bastos MT, Mercea D (2019) The Brexit botnet and user-generated hyperpartisan news. *Soc Sci Comput Rev* 37(1):38–54
- Beskow D (2020) Finding and characterizing information warfare campaigns. PhD thesis, Carnegie Mellon University

- Beskow DM, Carley KM (2018) Bot conversations are different: Leveraging network metrics for bot detection in Twitter. In: 2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM), IEEE, pp 825–832
- Beskow DM, Carley KM (2019a) Agent based simulation of bot disinformation maneuvers in Twitter. In: 2019 Winter Simulation Conference (WSC), IEEE, pp 750–761
- Beskow DM, Carley KM (2019b) Social cybersecurity: an emerging national security requirement. *Military Rev* 99(2):117
- Bradshaw S, Howard PN (2018) The global organization of social media disinformation campaigns. *J Int Affairs* 71(1.5):23–32
- Brin S, Page L (1998) The anatomy of a large-scale hypertextual web search engine. *Comput Netw ISDN Syst* 30(1–7):107–117
- Caldarelli G, De Nicola R, Del Vigna F, Petrocchi M, Saracco F (2020) The role of bot squads in the political propaganda on Twitter. *Commun Phys* 3(1):81. <https://doi.org/10.1038/s42005-020-0340-4>
- Carley KM (2020) Social cybersecurity: an emerging science. *Comput Math Org Theory* 26(4):365–381
- Carley KM, Columbus D, DeReno M, Reminga J, Moon IC (2008) ORA user's guide 2008. Carnegie Mellon University, Tech rep
- Carley KM, Cervone G, Agarwal N, Liu H (2018a) Social cyber-security. In: International Conference on Social Computing, Behavioral-Cultural Modeling and Prediction and Behavior Representation in Modeling and Simulation, Springer, pp 389–394
- Carley LR, Reminga J, Carley KM (2018b) ORA & NetMapper. In: International Conference on Social Computing, Behavioral-Cultural Modeling and Prediction and Behavior Representation in Modeling and Simulation, Springer
- Cha M, Gao W, Li CT (2020) Detecting fake news in social media: an Asia-Pacific perspective. *Commun ACM* 63(4):68–71
- Chong ZP (2015) #ge2015: Do you hear (what) the people sing? <https://medium.com/@thatswhathesaid/ge2015-do-you-hear-what-the-people-sing-a16cdf69fde>
- Cinelli M, Quattrociochi W, Galeazzi A, Valensise CM, Brugnoli E, Schmidt AL, Zola P, Zollo F, Scala A (2020) The COVID-19 social media infodemic. arXiv preprint [arXiv:200305004](https://arxiv.org/abs/200305004)
- Daryani M (2015) Silent majority: 6 lessons from Twitter data this #GE2015. <https://public.tableau.com/views/SilentMajorityGE2015/SilentMajority>
- Dorogovtsev SN, Goltsev AV, Mendes JFF (2006) K-core organization of complex networks. *Physical Review Letters* 96(4):040,601
- Dwyer T (2019) Media manipulation, fake news, and misinformation in the Asia-Pacific region. *J Contemp Eastern Asia* 18(2):9–15
- Ferrara E, Varol O, Davis C, Menczer F, Flammini A (2016) The rise of social bots. *Commun ACM* 59(7):96–104
- Ferrara E, Chang H, Chen E, Muric G, Patel J (2020a) Characterizing social media manipulation in the 2020 US presidential election. *First Monday*
- Ferrara E, Cresci S, Luceri L (2020b) Misinformation, manipulation, and abuse on social media in the era of COVID-19. *J Comput Soc Sci* 1–7
- Gastil J, Richards RC (2017) Embracing digital democracy: a call for building an online civic commons. *PS Polit Sci Polit* 50(3):758–763
- Geschke D, Lorenz J, Holtz P (2019) The triple-filter bubble: using agent-based modelling to test a meta-theoretical framework for the emergence of filter bubbles and echo chambers. *Br J Soc Psychol* 58(1):129–149
- Han K (2019) Big brother's regional ripple effect: Singapore's recent "fake news" law which gives ministers the right to ban content they do not like, may encourage other regimes in south-east asia to follow suit. *Index Censor* 48(2):67–69
- Humprecht E, Esser F, Van Aelst P (2020) Resilience to online disinformation: A framework for cross-national comparative research. *The International Journal of Press/Politics* p 1940161219900126
- James TS (2020) New development: Running elections during a pandemic. *Public Money Manag* 1–4
- Keller TR, Klinger U (2019) Social bots in election campaigns: theoretical, empirical, and methodological implications. *Polit Commun* 36(1):171–189
- King C, Bellutta D, Carley KM (2020) Lying about lying on social media: a case study of the 2019 Canadian Elections. In: International Conference on Social Computing, Behavioral-Cultural Modeling and Prediction and Behavior Representation in Modeling and Simulation, Springer, pp 75–85
- Krackhardt D, Stern RN (1988) Informal networks and organizational crises: an experimental simulation. *Soc Psychol Q* 51(2): 123–140

- Marcoux T, Obadimu A, Agarwal N (2020) Dynamics of online toxicity in the Asia-Pacific region. In: Pacific-Asia Conference on Knowledge Discovery and Data Mining, Springer, pp 80–87
- Mohar B (1989) Isoperimetric numbers of graphs. *J Comb Theory Ser B* 47(3):274–291
- Morstatter F, Pfeffer J, Liu H, Carley KM (2013) Is the sample good enough? Comparing data from Twitter's streaming API with Twitter's firehose. In: 7th International AAAI Conference on Weblogs and Social Media, ICWSM 2013, AAAI Press, pp 400–408
- Ng LHX, Yuan LJ (2020) Is this POFMA? Analysing public opinion and misinformation in a COVID-19 Telegram group chat. In: Workshop Proceedings of the 14th International AAAI Conference on Web and Social Media, vol 8, pp 1–8
- Noveck BS (2017) Five hacks for digital democracy. *Nature* 544(7650):287–289
- Pennebaker JW, Mehl MR, Niederhoffer KG (2003) Psychological aspects of natural language use: our words, our selves. *Annu Rev Psychol* 54(1):547–577
- Pérez-Rosas V, Kleinberg B, Lefevre A, Mihalcea R (2018) Automatic detection of fake news. In: Proceedings of the 27th International Conference on Computational Linguistics, pp 3391–3401
- Şen F, Wigand R, Agarwal N, Tokdemir S, Kasprzyk R (2016) Focal structures analysis: identifying influential sets of individuals in a social network. *Soc Netw Anal Min* 6(1):17
- Shallcross NJ (2017) Social media and information operations in the 21st century. *J Inf Warfare* 16(1):1–12
- Shao C, Hui PM, Wang L, Jiang X, Flammini A, Menczer F, Ciampaglia GL (2018) Anatomy of an online misinformation network. *PLoS ONE* 13(4):e0196087
- Shu K, Sliva A, Wang S, Tang J, Liu H (2017) Fake news detection on social media: a data mining perspective. *ACM SIGKDD Explor Newsl* 19(1):22–36
- Starbird K, Arif A, Wilson T (2019) Disinformation as collaborative work: Surfacing the participatory nature of strategic information operations. In: Proceedings of the ACM on Human-Computer Interaction 3(CSCW):1–26
- Tan N (2020) Electoral management of digital campaigns and disinformation in East and Southeast Asia. *Election Law J* 19(2)
- Tapscell R (2020a) Disinformation and cultural practice in Southeast Asia. In: Disinformation and Fake News, Springer, pp 91–101
- Tapscell R (2020b) Social media and elections in Southeast Asia: The emergence of subversive, underground campaigning. *Asian Stud Rev*, pp 1–18
- Tausczik YR, Pennebaker JW (2010) The psychological meaning of words: LIWC and computerized text analysis methods. *J Lang Soc Psychol* 29(1):24–54
- Traag VA, Waltman L, van Eck NJ (2019) From Louvain to Leiden: guaranteeing well-connected communities. *Sci Rep* 9(1):1–12
- Tucker JA, Guess A, Barberá P, Vaccari C, Siegel A, Sanovich S, Stukal D, Nyhan B (2018) Social media, political polarization, and political disinformation: a review of the scientific literature. Hewlett Foundation, Tech rep
- Udupa S (2018) Enterprise Hindutva and social media in urban India. *Contemp South Asia* 26(4):453–467
- Uyheng J, Carley KM (2020a) Bot impacts on public sentiment and community structures: Comparative analysis of three elections in the Asia-Pacific. In: International Conference on Social Computing, Behavioral-Cultural Modeling and Prediction and Behavior Representation in Modeling and Simulation. Springer
- Uyheng J, Carley KM (2020b) Bots and online hate during the COVID-19 pandemic: Case studies in the United States and the Philippines. *J Computat Soc Sci* 3:445–468
- Uyheng J, Carley KM (2021) Characterizing network dynamics of online hate communities around the COVID-19 pandemic. *Appl Netw Sci* 6(1):1–21
- Uyheng J, Magelinski T, Villa-Cox R, Sowa C, Carley KM (2019) Interoperable pipelines for social cyber-security: assessing Twitter information operations during NATO Trident Juncture 2018. *Computational and Mathematical Organization Theory*, pp 1–19
- Van Bavel JJ, Baicker K, Boggio PS, Capraro V, Cichocka A, Cikara M, Crockett MJ, Crum AJ, Douglas KM, Druckman JN, et al. (2020) Using social and behavioural science to support COVID-19 pandemic response. *Nat Hum Behav* pp 1–12
- Varol O, Ferrara E, Davis CB, Menczer F, Flammini A (2017) Online human-bot interactions: Detection, estimation, and characterization. In: Proceedings of the Eleventh International AAAI Conference on Web and Social Media (ICWSM 2017), pp 280–289
- Woolley SC (2016) Automating power: Social bot interference in global politics. *First Monday*
- Yee A (2017) Post-truth politics & fake news in Asia. *Glob Asia* 12(2):66–71

Zaini K (2020) Singapore in 2019: in holding pattern. *Southeast Asian Affairs* 2020(1):294–322
Gil de Zúñiga H, Veenstra A, Vraga E, Shah D (2010) Digital democracy: reimagining pathways to political participation. *J Inf Technol Polit* 7(1):36–51

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Joshua Uyheng is a PhD student in Societal Computing at Carnegie Mellon University. As a graduate researcher at the Center for Computational Analysis of Social and Organizational Systems (CASOS), he examines the network dynamics of online conflicts and digital disinformation, focusing on the Asia-Pacific and the Global South. He holds undergraduate degrees in mathematics and psychology from the Ateneo de Manila University in the Philippines.

Lynnette Hui Xian Ng is a PhD student in Societal Computing at Carnegie Mellon University. As a graduate researcher at the Center for Informed Democracy and Social Cybersecurity (IDeaS), she examines social cybersecurity and digital disinformation. She holds an undergraduate degree in computer science from the National University of Singapore.

Dr. Kathleen M. Carley is Professor of Computer Science and Director of the Center for Computational Analysis of Social and Organizational Systems (CASOS) at Carnegie Mellon University and the CEO of Netanomics.

Authors and Affiliations

Joshua Uyheng¹  · Lynnette Hui Xian Ng¹ · Kathleen M. Carley¹

✉ Joshua Uyheng
juyheng@cs.cmu.edu

Lynnette Hui Xian Ng
huixiann@cs.cmu.edu

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

¹ CASOS Center, Institute for Software Research, Carnegie Mellon University, 5000 Forbes Ave., Pittsburgh, PA 15213, USA