Characterizing Bot Networks on Twitter: An Empirical Analysis of Contentious Issues in the Asia-Pacific

Joshua Uyheng[1] [0000-0002-1631-6566] and Kathleen M. Carley[1] [0000-0002-6356-0238]

[1] CASOS, Institute for Software Research
Carnegie Mellon University, Pittsburgh PA, 15213
{juyheng, carley}@andrew.cmu.edu

Abstract. This paper empirically analyzes bot activity in contentious Twitter conversations using case studies from the Asia-Pacific. Bot activity is measured and characterized using a series of interoperable tools leveraging dynamic network analysis and machine learning. We apply this novel and flexible methodological framework to derive insights about information operations in three contexts: the senatorial elections in the Philippines, the presidential elections in Indonesia, and the relocation of a military base in Okinawa. Varying levels of bot prevalence and influence are identified across case studies. The presented findings demonstrate principles of social cyber-security in concrete settings and highlight conceptual and methodological issues to inform further development of analytic pipelines in studying online information operations.

Keywords: Information Operations, Social Cyber-Security, Asia-Pacific.

1 Introduction

Online social networks (OSNs) have provided citizens worldwide with a powerful platform to participate in political discourse and exchange views at an unprecedented scale [1, 2]. However, targeted information operations on these platforms curb such democratizing potentials by introducing artificial influences into online conversations [3]. Utilizing artificial agents like automated bots, information operations exploit the network structure of the social media platform (e.g., who speaks to whom) to manipulate dominant topics of discussion (e.g., who speaks about what). As political discourse increasingly encompasses the digital sphere, it is of prime significance for researchers and policymakers to develop integrated and effective frameworks for detecting and characterizing information operations across a variety of contexts [4].

ACKNOWLEDGMENTS. This work is supported in part by the Office of Naval Research under the Multidisciplinary University Research Initiatives (MURI) Program award number N000141712675 Near Real Time Assessment of Emergent Complex Systems of Confederates, BotHunter award number N000141812108, award number N00014182106 Group Polarization in Social Media: An Effective Network Approach to Communicative Reach and Disinformation, and award number N000141712605 Developing Novel Socio-computational Methodologies to Analyze Multimedia-based Cyber Propaganda Campaigns. This work is also supported by the center for Computational Analysis of Social and Organizational Systems (CASOS). 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 ONR or the U.S. government.

This paper proposes a novel and flexible methodological framework for understanding bot-driven information operations on Twitter. In the context of three contentious issues in the Asia-Pacific, we utilize a series of interoperable tools leveraging machine learning and dynamic network analysis to (a) detect user accounts that exhibit bot-like characteristics, (b) examine the topics of discussion in which such agents are involved, and (c) assess their overall impact on the online conversation [5, 6]. On a basic level, our findings contribute valuable insights into information operations in a highly significant geopolitical region. More broadly, we illustrate the value of interoperable pipelines for analyzing bot activity from the perspective of social cybersecurity [4, 7]. Finally, we provide practical considerations for stakeholders involved in the cases examined, especially as regards mapping artificial influences on contentious political discourse.

1.1 Information Operations in Online Social Networks

Over the past decade, social media sites like Facebook and Twitter have become instrumental in facilitating online communication in both personal and large-scale contexts. In the political sphere, OSNs have introduced a highly democratized space for public discourse. The open lines of communication introduced by OSNs allow for diverse actors within and across societies to exchange ideas and information in a relatively unconstrained manner [2]. However, for the same reasons, OSNs are also susceptible to manipulation. In the present work, we are interested in empirically detecting and characterizing forms of manipulation such as targeted information operations.

The use of information operations to influence public opinion is not new. However, targeted communications in the context of OSNs acquire a vastly accelerated potential to spread malignant messages. Disinformation, polarization, and even terrorist radicalization have all been associated with the unregulated proliferation of OSNs [8–10]. Extensive work shows how OSNs can be used to heighten intergroup conflict [11], recruit support for extremist organizations [12], or nudge electoral decisions [13]. By combining human and automated agents (i.e., bots), information operations leverage the dynamics of OSNs to affect real-world outcomes. Key stakeholders in such operations may span individuals and organizations to nation-states.

1.2 Social Cyber-Security with Interoperable Tools

Scholarship in the burgeoning field of social cyber-security seeks to understand and combat digital technologies utilized for adversarial purposes using a ‘multidisciplinary and ‘multimethodological’ approach [4]. In this paper, we propose three research questions which operationalize foundational facets of social cyber-security inquiry. What is the level of bot activity in the online conversation? What messages are most targeted by bot activity? How does bot activity influence the social network? As we aim to show, these questions offer a general yet insightful framework for characterizing information operations across a variety of contentious political issues.

We further propose a novel and flexible methodological pipeline of interoperable tools to empirically address these questions in concrete contexts. The tools embedded in our framework leverage various concepts and techniques in network science and

machine learning. Network science studies how entities and their relationships can be modeled using a graph structure with nodes and edges [14]. Dynamic network analysis characterizes large, high-dimensional, and time-variant graphs to quantify the structure of online conversations, the importance of some actors over others, and the relationships between agents and the topics they discuss. Several machine learning algorithms are also embedded in our pipeline. First, Bothunter deploys individual- and network-level features to compute a probability that a given user is a bot [6]. The algorithm is a random forest trained on a labeled dataset of known bots. The second tool identifies whether a user is a news agency, a government account, a company, or a celebrity, among other pre-identified account types known to drive significant Twitter traffic. This tool is based on a neural network model trained on a labeled dataset of user-provided account descriptions and their latest 20 tweets.

At this juncture, we note that while these concepts are not new in the study of OSNs and bot activity, the key feature of our proposed framework is its integrated pipeline of analysis. Whereas numerous studies have produced novel means of analyzing social media data and bot behavior online, few demonstrate frameworks for synergistically and effectively using these tools to distill holistic and actionable findings. In a practical setting, diverse data types and research questions must be examined concurrently, and therefore ought to be compatible for the generation and triangulation of insights. As our findings illustrate in subsequent sections, this paper demonstrates an informative model of such an integrated framework.

1.3 Contentious Issues in the Asia-Pacific

To illustrate the utility of our proposed pipeline, we examine three case studies of contentious issues in the Asia-Pacific. We collect Twitter conversations surrounding the Philippine senatorial elections, the Indonesian presidential elections, and campaigns against the relocation of a military base to Oura Bay in Okinawa. In all cases, tweets were in English, the local language (e.g., Filipino), or a blend of both.

#Halalan2019: The Philippine Senatorial Elections. The Philippines holds midterm elections after three years in the sitting President's six-year term. In May 2019, midterm elections will decide twelve senatorial positions, comprising half of the seats of the nationally elected Senate [15]. Set against the backdrop of the polarizing regime of President Rodrigo Duterte, the senatorial race constitutes a high-stakes contest for legislative power in the country. The race features major political figures representing administration-backed candidates and a coalition of opposition candidates.

#Pemilu2019: The Indonesian Presidential Elections. In April 2019, Indonesia will hold general elections to fill a vast majority of its high-ranking political posts [16]. Foremost among these contests will be the race for the presidency, for which incumbent Joko Widodo will reiterate his 2014 clash against former general Prabowo Subianto. Whereas substantial controversy has surrounded Widodo's term for its economic underperformance and unconvincing stance on China, Subianto's "Make Indonesia Great Again" rhetoric has likewise been subjected to scrutiny for its similarity to

Donald Trump's populist rhetoric during the 2016 US presidential elections. Both candidates will campaign for a five-year term of the nation's leadership.

#StandWithOkinawa: Military Base Relocation in Oura Bay. The final case concerns the relocation of a US military base to Oura Bay in the Okinawa prefecture [17]. While relocation plans had been in progress since the mid-1990s, resistance from the local population delayed construction until a Supreme Court ruling in 2016. Concerns of the opposition include people's safety in areas linked to military facilities, the purportedly dubious proceedings legalizing the relocation's approval, and environmental effects of the construction which may impact dugong habitats.

2 Data and Methods

To analyze information operations in the three Asia-Pacific contexts described, this study employed an integrated methodological framework. Utilizing interoperable network science and machine learning tools, we characterized the users, topics, and overall network influence associated with bot-like accounts on public Twitter data harvested with the STREAM and REST APIs.

2.1 Datasets

For each case study, a set of initial hashtags was generated based on mainstream media coverage of the issues. From this initial set, further search terms were iteratively identified and incorporated into the search. For the Philippine and Indonesian elections, search terms consisted of official election hashtags (e.g., #Halalan2019 in the Philippines; #Pemilu2019 in Indonesia) as well as the names and campaign slogans of candidates running for office. Meanwhile, for the case in Japan, the initial hashtag was #standwithokinawa, which encapsulated the campaign to resist military base relocation. From this initial search term, new terms were added including hashtags in Japanese. Table 1 summarizes the data collection methods used for each case study, the inclusive dates of the tweets analyzed, and the total number of tweets examined.

Table 1. Description of case study datasets.

Case Study	Sample Hashtags	Inclusive Dates	Tweets Collected
Philippine Elections	#Halalan2019	Oct 2018 – Jan 2019	473993
Indonesian Elections	#Pemilu2019	Dec 2018 – Jan 2019	305959
Stand with Okinawa	#standwithokinawa	Dec 2018 – Jan 2019	260562

2.2 Analytic Procedure

The overall analytic scheme followed a sequence of three broad stages: (a) user analysis, (b) topic analysis, and (c) network analysis. Implementation flow is summarized in Figure 1.

During user analysis, Bothunter and role identity algorithms were deployed in parallel. The results of both algorithms were converged to construct a contingency table of predicted bots and possible role identities. The latter analysis clarified bot predictions by providing an alternative explanation for the 'bot-like' behavior identified by the former tool. Only bots concurrently classified as non-special actors by the role identity algorithm were retained, providing conservative estimates of bot activity.

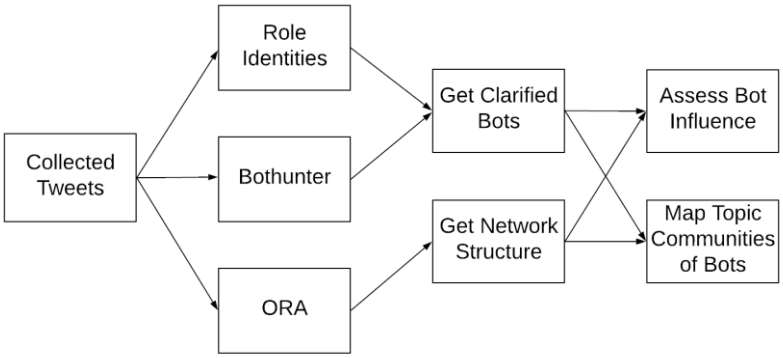

Fig. 1. Interoperable pipeline of social cyber-security tools.

Topic communities in the Twitter conversation for each case study were concurrently identified using ORA. ORA visualized Twitter users as nodes connected by weighted edges depending on the degree to which users had cited the same hashtags. Using Louvain clustering, we identified unique topics of discussion based on the hashtags used by the identified topic communities. The Louvain clustering algorithm optimizes modularity [18], which measures the network density within candidate clusters relative to vertices outside these clusters. Converged with the identified bots, we also identified which topics featured the highest level of predicted bot activity.

Finally, network analysis on ORA examined the higher-level Twitter interactions which took place in each online conversation. Simultaneously computing a host of network centrality metrics, ORA generated reports of Twitter users identified as 'super spreaders' and 'super friends' who wielded substantial influence in the collected dataset through their high-ranking centrality scores on communication networks (i.e., their messages are widely spread) or friendship networks (i.e., they are friends with many users), respectively. Such agents play an instrumental role in the online discussion. Interpreted in conjunction with the user analysis and topic analysis, this high-level approach provided a rich map of which messages (i.e., topics) are being propagated by whom (i.e., bots), thereby assessing the impact of potential information operations on the overall discussion (i.e., network influencers).

3 Results

Detected bots represented a non-negligible proportion of users across social networks, with notable country-level differences in general bot activity. Bot operations further

varied across contexts, with bots attempting to drive intergroup opposition, rally intragroup support, or propel various facets of a contentious topic.

3.1 User Analysis: Bots and Special Actors

Bothunter produced probabilities that each user in the datasets was an automated bot. A 60% probability threshold was used for positive cases based on previous validation procedures with Bothunter. The role identity algorithm identified whether users belonged to special classes of actors with characteristically numerous tweets and followers. Table 2 cross-tabulates these results. Numbers in bold indicate the final bot estimates given that the role identity algorithm did not indicate they belonged to special actor classes. Percentages are relative to the total number of users in each dataset.

Table 2. Detected bots with predicted role identities.

Case Study	Role Identities	Unique Users	Detected Bots
Philippine Elections	Normal	63260 (78.79%)	**9073 (11.30%)**
	Government	7424 (9.25%)	1189 (1.48%)
	News Agency	3538 (4.41%)	510 (0.64%)
	News Reporter	2950 (3.67%)	630 (0.78%)
	Company	802 (1.00%)	103 (0.13%)
	Celebrity	2015 (2.51%)	172 (0.21%)
	Sports	300 (0.36%)	44 (0.05%)
Indonesian Elections	Normal	21987 (87.48%)	**2568 (10.22%)**
	Government	722 (2.87%)	95 (0.28%)
	News Agency	471 (1.87%)	26 (0.10%)
	News Reporter	60 (0.24%)	10 (0.04%)
	Company	1569 (6.24%)	215 (0.86%)
	Celebrity	230 (0.92%)	32 (0.13%)
	Sports	95 (0.28%)	6 (0.02%)
Stand with Okinawa	Normal	31428 (98.09%)	**8750 (27.31%)**
	Government	103 (0.32%)	30 (0.09%)
	News Agency	171 (0.53%)	22 (0.07%)
	News Reporter	23 (0.08%)	10 (0.03%)
	Company	261 (0.81%)	55 (0.17%)
	Celebrity	25 (0.08%)	3 (0.01%)
	Sports	29 (0.09%)	2 (0.01%)

A non-negligible proportion of captured users across the three datasets were detected as bots by our algorithms. Across the contentious issues analyzed, between 10% and nearly 30% of all captured Twitter users were bots or bot-like, second only to ordinary accounts in the datasets but outnumbering news agencies and government ac-

counts. Country-level bot statistics appeared to be markedly different, with Japan having the highest number of bots and Indonesia the least among the case studies.

3.2 Topic Analysis: Bot-Active Hashtags

Topics in the Twitter conversations were identified by examining the most frequently used hashtags of Louvain clusters of agents. Top hashtags for the topic groups having the highest number of bots are presented in Table 3. Bot percentages are given relative to the total number of users in each topic group. Labels for topics were interpreted based on the top hashtags validated by manually reading sample tweets.

Table 3. Top hashtags in topics with most bot activity.

Case Study	Topic	Bot Activity	Hashtags
Philippine Elections	Opposition campaign (focused on Roxas)	22.93%	OtsoDiretso, MarRoxas, RoxasTayo, OposisyonKoalisyon, TheLeaderIWant
	Opposition campaign (focused on Hilbay)	12.82%	HilbayForSenator, MarRoxas, TeamPhilippines, OtsoDiretso, OposisyonKoalisyon
	Anti-opposition, pro-Duterte	13.18%	NoToLiberalParty2019, NeverAgainToLP, ImeeSolusyon, NoToNeri, TeamDDS2019
Indonesian Elections	Campaign for peaceful elections	11.11%	pilpres, pemiludamai, KPU, Pilpres2019, kampanye
	Campaign for opposition (Subianto)	6.60%	AkalSehatPilihPrabowoSandi, SumbarMemilihPrabowoSandi, 2019PrabowoPresidentRI
	Campaign for incumbent (Widodo)	1.40%	OrangBaikDukungJokowi, 01Indonesiamaju, jokowilagi, 2019TetapJokowi
Stand with Okinawa	White House petition against base	25.77%	辺野古の海を埋め立てないで, ホワイトハウスへ36万筆署名, Democracy
	Protesting Shinzo Abe	24.18%	ヤバすぎる安倍政権, ヤバすぎる緊急事態条項
	Referendum against military landfill	22.93%	県民投票, 辺野古県民投票, 普天間移設問題

For each topic above, the main hashtags (e.g., #Halalan2019, #standwithokinawa) were omitted as they appeared in all topics. In elections, both administration and opposition candidates were discussed by bot-like accounts, suggesting that information operations may be used by both parties, or that bots aim to influence messaging about all candidates. Further research may probe this question. In the Okinawa case, it appears that both local and international issues are amplified by such accounts.

3.3 Network Analysis: Influencer Assessment

Finally, top influencers were identified by network centrality. Super spreaders and super friends are identified in Table 4 with Bothunter scores. Asterisks mark verified accounts. Italics are for accounts suspended or moved since data collection. Bolded names may require deeper analysis due to high influence and bot-like behavior.

Table 4. Super spreaders and super friends.

Case Study	Super Spreaders	Super Friends
Philippine Elections	rapplerdotcom* (0.37)	ru6dy9 (0.50)
	MARoxas* (0.41)	raincyrainy (0.52)
	ATajum (0.44)	*mariagarciaah (0.42)*
	cnnphilippines* (0.20)	*MayDPoresBeWidU* (0.54)
	mariagarciaah (0.42)	MelLegaspi1 (0.44)
	AsecMargauxUson (0.51)	*AsecMargauxUson (0.51)*
	ru6dy9 (0.50)	jvejercito* (0.69)
	BembangBiik (0.37)	**BoyoKiss (0.65)**
Indonesian elections	Sandiuno* (0.47)	Addarul1 (0.57)
	CakKhum (0.25)	HotPepperminTea (0.51)
	Gerindra* (0.53)	abiid_d (0.56)
	Addarul1 (0.57)	abiyyu231299 (0.38)
	02Sandiaga (0.45)	Rusydi_riau40 (0.52)
		MangajatsCkp (0.55)
		Bagusalghazali (0.54)
Stand with Okinawa	**surumegesogeso (0.65)**	tkatsumi06j (0.30)
	robkajiwara (0.57)	affluencekana (0.46)
	ISOKO_MOCHIZUKI (0.65)	**sabor_sabole (0.74)**
	times_henoko (0.35)	robkajiwara (0.57)
	29ryukyu (0.33)	HempHere (0.34)
	mr_naha_das (0.49)	29_momechabo (0.57)
	BFJNews* (0.59)	HIROMI150303 (0.48)
		ActSludge (0.50)

In each case, major influencers appear to be a mix of politicians (e.g., MARoxas, Sandiuno), media outlets (e.g., rapplerdotcom, BFJNews), and bot-like agents (e.g., BoyoKiss, sabor_sabole). While the first two categories are expected, the last are noteworthy especially as they dominate super friends. While their messages are not entirely influential (e.g., retweeted), their vast connections (e.g., followers) might still

make their messages visible to many. The present analysis thus flags accounts which may require deeper investigation of their activities and impact on public discourse.

4 Discussion

This paper analyzed three case studies of contentious Twitter conversations in the Asia-Pacific. Bots and bot-like accounts appeared to be non-negligible in terms of prevalence, participating in diverse activities such as promoting causes, antagonizing the opposition, and elaborating on various facets of an issue [7]. Country-level variations were observed, suggesting that information operations are more widely utilized in certain nations over others. Nonetheless, in each case, bots and bot-like accounts wielded significant influence, boasting high centrality measures on par with media outlets and government officials. Accounts requiring deeper investigation were flagged, among which some had already been suspended or deactivated, potentially providing partial validation of our predictions.

In the growing field of social cyber-security, our results illustrate the utility and need for integrated approaches to assess, characterize, and combat targeted information operations online. Pipelined interoperable tools triangulate insights beyond a single analytical approach. By providing a high-level map of potentially suspicious activity, our findings enable more in-depth inquiry into specific operations informed by specialized knowledge of local sociopolitical dynamics. Due to the ongoing nature of each campaign upon the writing of this paper, we note that new developments may arise beyond the scope of the present analysis. Within and beyond the Asia-Pacific, potential for covert intervention calls for rigorous vigilance from a diversity of fronts. We aimed to show how such vigilance might be implemented in practice.

References

1. Kahne, J., Bowyer, B.: The Political Significance of Social Media Activity and Social Networks. Political Communication. 35, 470–493 (2018). https://doi.org/10.1080/10584609.2018.1426662.
2. McGarty, C., Thomas, E.F., Lala, G., Smith, L.G.E., Bliuc, A.-M.: New Technologies, New Identities, and the Growth of Mass Opposition in the Arab Spring. Political Psychology. 35, 725–740 (2014). https://doi.org/10.1111/pops.12060.
3. Bandeli, K.K., Agarwal, N.: Analyzing the role of media orchestration in conducting disinformation campaigns on blogs. Comput Math Organ Theory. (2018). https://doi.org/10.1007/s10588-018-09288-9.
4. Carley, K.M., Cervone, G., Agarwal, N., Liu, H.: Social cyber-security. In: Social, Cultural, and Behavioral Modeling - 11th International Conference, SBP-BRiMS 2018, Proceedings. pp. 389–394. Springer Verlag (2018). https://doi.org/10.1007/978-3-319-93372-6_42.
5. Benigni, M., Joseph, K., Carley, K.M.: Mining online communities to inform strategic messaging: practical methods to identify community-level insights. Comput

Math Organ Theory. 24, 224–242 (2018). https://doi.org/10.1007/s10588-017-9255-3.
6. Beskow, D.M., Carley, K.M.: Its all in a name: detecting and labeling bots by their name. Comput Math Organ Theory. (2018). https://doi.org/10.1007/s10588-018-09290-1.
7. Beskow, D.M., Carley, K.M.: Social cybersecurity: An emerging national security requirement. Military Review. March-April, (2019).
8. Bail, C.A., Argyle, L.P., Brown, T.W., Bumpus, J.P., Chen, H., Hunzaker, M.B.F., Lee, J., Mann, M., Merhout, F., Volfovsky, A.: Exposure to opposing views on social media can increase political polarization. Proceedings of the National Academy of Sciences. 115, 9216–9221 (2018). https://doi.org/10.1073/pnas.1804840115.
9. Tucker, J.A., Guess, A., Barberá, P., Vaccari, C., Siegel, A., Sanovich, S., Stukal, D., Nyhan, B.: Social media, political polarization, and political disinformation: A review of the scientific literature. Hewlett Foundation (2018).
10. Ong, J.C., Cabanes, J.V.A.: Architects of networked disinformation: Behind the scenes of troll accounts and fake news production in the Philippines. Newton Tech4Dev Network (2018).
11. Babcock, M., Cox, R.A.V., Kumar, S.: Diffusion of pro- and anti-false information tweets: the Black Panther movie case. Comput Math Organ Theory. (2018). https://doi.org/10.1007/s10588-018-09286-x.
12. Klausen, J.: Tweeting the *Jihad* : Social Media Networks of Western Foreign Fighters in Syria and Iraq. Studies in Conflict & Terrorism. 38, 1–22 (2015). https://doi.org/10.1080/1057610X.2014.974948.
13. Allcott, H., Gentzkow, M.: Social Media and Fake News in the 2016 Election. Journal of Economic Perspectives. 31, 211–236 (2017). https://doi.org/10.1257/jep.31.2.211.
14. Al-Garadi, M.A., Varathan, K.D., Ravana, S.D., Ahmed, E., Mujtaba, G., Khan, M.U.S., Khan, S.U.: Analysis of online social network connections for identification of influential users: Survey and open research issues. ACM Computing Surveys. 51, 16:1-16:37 (2018).
15. Bueza, M.: Survey says: How 2019 senatorial bets are faring so far, https://www.rappler.com/newsbreak/iq/220707-senatorial-candidates-survey-performance-2019-elections, (2019).
16. Cochrane, J.: Indonesia's presidential race takes shape, in shadow of hard-line Islam, https://www.nytimes.com/2018/08/11/world/asia/indonesia-presidential-election.html, (2018).
17. Japan begins filling in Henoko Bay in Okinawa to make room for unpopular US base, https://www.japantimes.co.jp/news/2018/12/14/national/japan-starts-landfill-work-move-unpopular-u-s-base-okinawa, (2018).
18. Blondel, V.D., Guillaume, J.-L., Lambiotte, R., Lefebvre, E.: Fast unfolding of communities in large networks. Journal of Statistical Mechanics: Theory and Experiment. 2008, P10008 (2008). https://doi.org/10.1088/1742-5468/2008/10/P10008.