Unsupervised Characterization of State-Sponsored Twitter Information Operations with Multi-view Clustering?

Joshua Uyheng[1] [0000 0002 1631 6566], Iain Cruickshank[1] [0000 0002 4205 5806], and Kathleen M. Carley[1] [0000 0002 6356 0238]

CASOS Center, Institute for Software Research, Carnegie Mellon University, Pittsburgh PA 15213, USA
fjuyheng, icruicks, kathleen.carleyg@cs.cmu.edu

Coordinated actors use information maneuvers to shift public narratives and alter the information ow between individuals and groups [4]. The impacts of such campaigns reverberate not only through cyberspace, but also in high-stakes,

users or just the text being emitted by users. As such, if a user does not engage in a certain type of interaction or post any text, they cannot be e ectively clustered by that one source of information. So, using more views of the users can provide more information to not only provide better clusters, but also account for things like partially complete views of the data [11].

Through the use of *multi-view clustering*, we introduce an unsupervised method to describe the functional elements of state-sponsored information operations. These data-driven, functional characterizations of actors help researchers better understand how disinformation operators organize tasks and undertake disinformation campaigns. This potentially unmasks novel tactics in information maneuvers, without relying on extensive and expensive labeled data [7]. This information is vital in helping platforms and governments to better identify and counter disinformation campaigns as they emerge against the intersecting volatile contexts of cyberspace and geopolitics [4].

We use a newly developed technique called Multi-view Modularity Clustering (MVMC) to cluster all of the users identi ed as part of a nation-state's information operation [6]. This approach identi es functional components of users based on both the text they emit and their patterns of interaction. To perform the multi-view clustering, we form text, hashtag, retweet, and interaction views out of Twitter data from all of the users. By creating graphs from the non-network views, the MVMC method simultaneously clusters all of the view graphs using an iterative optimization on an adapted version of network modularity [10]. The end result is a set of labels for all of the users. Using these labels, we further characterize each cluster based on the BEND framework of information maneuvers. We use the ORA software to empirically measure cluster characteristics associated with: (a) network maneuvers, which in uence the ow of information by altering who talks to whom; and (b) narrative maneuvers, which shape positive and negative storylines regarding objects of online discussion [2].

We run this analysis using the Information Operations dataset released by Twitter on China (May 2020). Text features are extracted using the Netmapper software [5], which relies on multilingual lexicons to extract several psycholinguistic measures of interest in information operations [8]. These include the use of abusive words, di erent pronoun types, emojis, and identity terms, among others. These text features are formed into view graph by using a symmetric k-Nearest Neighbor graph learning procedure. Interactions, on the other hand, are derived using patterns of mentions, replies, quotes, and retweets in the dataset. Interaction networks are produced by representing users in the data as nodes in a graph, with the number of times they communicate corresponding to the edge weights for each interaction type. Upon observing a distinct increase in the tweets associated with this dataset after November 2019, we produced two sets of view graphs: one for before, and one for after this uptick.

Table 1 summarizes the BEND maneuvers detected through the proposed MVMC-driven framework. Across both time periods, accounts dealing with the Hong Kong protests employed a combination of negative storylines to paint protestors in a bad light, while invoking positive storylines about China and

Table 1. Summary of BEND maneuvers measured on MVMC clusters of state-sponsored actors.

MVMC Group	Storylines	BEND Maneuvers
First Era		
Hong Kong Protests	Target protestors of repressive Hong Kong law	Back, Enhance, Neutralize, Dismay
	Respond to violence with iron st	Enhance, Neutralize, Dismay
	Target pro-democracy electoral candidate	Back, Excite
	Love Hong Kong by supporting police	Back, Enhance
Guo Wen Gui	Link GWG to Steve Bannon and rapists	Neutralize, Dismay
	Link GWG to far-right movements	Enhance, Dismiss
Covid-19	Link Wuhan situation to need for police	Back, Enhance, Neutralize, Dismay
Second Era		

as explain maneuvers to introduce new framing. Finally, for both time periods, agents discussing COVID-19 were concentrated in one cluster, indicating highly organized behavior. Furthermore, while accounts which existed prior to November 2019 sought to link COVID-19 to the importance of police, post-November accounts sought to pin COVID-19 responsibility to the US President.

Through the use of multi-view clustering, we introduce an unsupervised method to describe the functional elements of state-sponsored information operations [4,9]. These data-driven, functional characterizations of actors help researchers better understand how disinformation operators organize tasks and undertake disinformation campaigns. This potentially unmasks novel tactics in information maneuvers, without relying on extensive and expensive labeled data. This information is vital in helping platforms and governments to better identify and counter disinformation campaigns as they emerge against the intersecting volatile contexts of cyberspace and geopolitics [3].

References

1. Beskow, D.M., Carley, K.M.: Bot conversations are di erent: Leveraging network metrics for bot detection in twitter. In: 2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM). pp. 825{832. IEEE (2018)
2. Beskow, D.M., Carley, K.M.: Social cybersecurity: An emerging national security requirement. Military Review **99**(2), 117 (2019)
3. Bradshaw, S., Howard, P.N.: The global organization of social media disinformation campaigns. Journal of International A airs **71**(1.5), 23{32 (2018)
4. Carley, K.M., Cervone, G., Agarwal, N., Liu, H.: Social cyber-security. In: International Conference on Social Computing, Behavioral-Cultural Modeling and Prediction and Behavior Representation in Modeling and Simulation. pp. 389{394. Springer (2018)
5. Carley, L.R., Reminga, J., Carley, K.M.: ORA & NetMapper. In: International Conference on Social Computing, Behavioral-Cultural Modeling and Prediction and Behavior Representation in Modeling and Simulation. Springer (2018)