

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

Joshua Uyheng¹[0000–0002–1631–6566], Iain Cruickshank¹[0000–0002–4205–5806],
and Kathleen M. Carley¹[0000–0002–6356–0238]

CASOS Center, Institute for Software Research, Carnegie Mellon University,
Pittsburgh PA 15213, USA
{juyheng, icruicks, kathleen.carley}@cs.cmu.edu

Coordinated actors use information maneuvers to shift public narratives and alter the information flow between individuals and groups [4]. The impacts of such campaigns reverberate not only through cyberspace, but also in high-stakes, real-world settings [3]. Using previously labeled data, machine learning tools abound for automatically identifying potential misinformation, rumors, bot accounts, hate speech, and other elements which may attend attempts to manipulate public discourse [1]. Network science likewise provides systematic frameworks for representing complex patterns of online communication. This facilitates the principled identification of actors holding various types of influence over the conversation, acting as hubs for information, bridges between groups, or trying to break groups apart [2].

However, many frameworks developed focus on specific components of information operations in isolation [4]. It is unclear how a larger information operations group may use all of these different approaches as part of a campaign. As information operations unfold in real time, they achieve their objectives in multiple, dynamic ways. Further, it may not be known in advance which analyses are most relevant to which facets of the campaign. Hence, while existing models – especially those founded on supervised machine learning algorithms – may generate rich and accurate predictions on precise dimensions of an information operation, they do not account for the multifaceted and novel nature of influence campaigns as a whole [9].

Multi-view clustering offers an unsupervised methodology for characterizing information operations by finding groups of users by simultaneously considering their attributes as well as their interactions [6]. Common approaches to clustering are hampered by their exclusive focus on just networks of interactions between

* 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.

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 effectively 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 unmask 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 identified as part of a nation-state’s information operation [6]. This approach identifies 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 influence the flow 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, different 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
<i>First Era</i>		
Hong Kong Protests	Target protestors of repressive Hong Kong law	Back, Enhance, Neutralize, Dismay
	Respond to violence with iron fist	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
<i>Second Era</i>		
Hong Kong Protests	Comparing protestors to thugs against justice	Enhance, Neutralize, Dismay
	Protect Hong Kong by supporting the police	Back, Enhance, Neutralize, Dismay
	HK is part of China	Excite
	Ridiculing Hong Kong elections	Enhance, Neutralize, Dismay
Guo Wen Gui	Calling GWG a fugitive	Back, Explain, Enhance, Neutralize, Distort
	Calling GWG and western media hypocrites	Explain
	Linking GWG to right-wing operatives	Excite
Covid-19	Linking Covid-19 to Trump	Enhance, Neutralize, Dismay

the police. This is exemplified through narrative maneuvers like dismay, which links concepts to negative-valence terms; versus positive maneuvers like enhance, which encourages further conversation through community-related words, and excite, which associates concepts with positive-valence terms. Network maneuvers included backing opinion leaders in the cluster, and neutralizing others' opinions, both with systematic retweeting behavior. Similar maneuvers were implemented to frame Guo Wen Gui as a criminal and his supporters in the West as hypocrites, with the addition of distort maneuvers which confuse a topic, as well

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 2019 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 different: 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 Affairs* **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)
6. Cruickshank, I.J.: Multi-view clustering of social-based data. Ph.D. thesis, PhD thesis, Carnegie Mellon University (2020)
7. Cruickshank, I.J., Carley, K.M.: Characterizing communities of hashtag usage on twitter during the 2020 covid-19 pandemic by multi-view clustering. *Applied Network Science* **5**(1), 1–40 (2020)
8. Tausczik, Y.R., Pennebaker, J.W.: The psychological meaning of words: LIWC and computerized text analysis methods. *Journal of Language and Social Psychology* **29**(1), 24–54 (2010)
9. Uyheng, J., Magelinski, T., Villa-Cox, R., Sowa, C., Carley, K.M.: Interoperable pipelines for social cyber-security: Assessing Twitter information operations during NATO Trident Juncture 2018. *Computational and Mathematical Organization Theory* pp. 1–19 (2019)
10. Wang, H., Yang, Y., Liu, B., Fujita, H.: A study of graph-based system for multi-view clustering. *Knowledge-Based Systems* **163**, 1009–1019 (2019)
11. Yang, Y., Wang, H.: Multi-view clustering: A survey. *Big Data Mining and Analytics* **1**(2), 83–107 (2018)