

Coordinated through a Web of Images: Analysis of Image-based Influence Operations from China, Iran, Russia, and Venezuela

Lynnette Hui Xian Ng¹, J.D. Moffitt¹, Kathleen M. Carley¹

¹School of Computer Science, Carnegie Mellon University
5000 Forbes Ave
Pittsburgh, Pennsylvania 15213
{huixiann, jdmoffitt, kathleen.carley}@cs.cmu.edu

Abstract

State-sponsored online influence operations typically consist of coordinated accounts exploiting the online space to influence public opinion. Accounts associated with these operations use images and memes as part of their content generation and dissemination strategy to increase the effectiveness and engagement of the content. In this paper, we present a study of images from the PhoMemes 2022 Challenge originating from the countries China, Iran, Russia, and Venezuela. First, we analyze the coordination of images within and across each country by quantifying image similarity. Then, we construct Image-Image networks and image clusters to identify key themes in the image influence operations. We derive the corresponding Account-Account networks to visualize the interaction between participating accounts within each country. Finally, we interpret the image content and network structure in the broader context of the organization and structure of influence operations in each country.

Introduction

The global information environment connects the physical, information, and cognitive domains. Engagements in the physical and information domains shape people’s world attitudes, beliefs, and world views in the cognitive domain. The rise of the internet age and the development of social media platforms provides state and non-state actors a unique opportunity to exploit the information environment to conduct influence operations (Stout 2017).

Influence operations can be broadly defined as coordinated, integrated and synchronized efforts to accent communications to affect attitudes and behaviors (Larson et al. 2009). Several frameworks have been proposed to characterize influence operations: a process of defining broad objectives then zooming into detailed strategies and deploying negotiation techniques (Larson et al. 2009); fusing internet and military concepts in a kill-chain approach Bergh (2020); mapping pathways of information on social media platforms (Ng and Taeihagh 2021), and characterizing narrative and network maneuvers (Carley 2020).

Studies have indicated that China, Iran, Russia, and Venezuela have substantial state capacity and resources dedicated to conducting influence operations (Bradshaw and

Howard 2019). In addition, these countries’ coordinated influence operations can be categorized in terms of their content and communication features (Alizadeh et al. 2020). Prior work identified Twitter posts from Russia as spreading propaganda and politically biased information during the 2016 US election (Badawy et al. 2019). There is also strong evidence of automated social media manipulation operations stemming from Iran (Bradshaw and Howard 2017).

State-sponsored accounts do not rely solely on textual content but also leverage the expressive power of images, e.g., using politically and ideologically-charged memes (Rowett 2018). The use of images contributes to increasing message engagement (Wang et al. 2020), which can influence online political opinion (Schill 2012). While there is a long stream of work analyzing texts used in influence operations, image analysis in the context of influence operations has been less studied.

In this paper, we begin filling this gap by analyzing the use of images by state-sponsored accounts in information diffusion. We adopt network-based approaches in analyzing and comparing the coordination structures among the four countries. This approach draws connections between images to indicate their similarity, then further draws connections between accounts when they both coordinate via sharing a similar image. The resultant networks allow visual inspection and comparison of the image and account coordination structure that describes the influence operation of each country. Implementing an unsupervised image comparison and clustering approach provides more flexibility than pre-defined categories in a supervised setting. It is capable of adapting as images and memes evolve during the course of an influence operation. We not only examine the image and account coordination structure through a computational pipeline but also contextualize the observations with the influence operation capabilities of each country as a backdrop.

Contributions. Using images obtained from identified Twitter inauthentic accounts from four countries, we analyze the image-based influence operations by country and collectively. In this study, we make the following contributions:

1. We develop a systematic methodology for analyzing coordinated image dissemination as part of an influence operation strategy. We analyze the use of images in influence operations in four countries and cross-country image coordination through image similarity networks.

2. We compare the coordination strategy of accounts in image information dissemination through Account-Account networks across the four countries.
3. We provide an overview of the image messaging strategy of each country in the broader context using studies on the structure and organization of influence operations in each country.

Related Work

Influence Operations. Social media is an effective tool for influencing individuals’ opinions (Bakshy et al. 2011). The coordinated operation to influence the opinions of large groups of people can threaten the social fabric, as witnessed in the spread of #StopTheSteal narratives during the 2021 Capitol Riots (Ng, Cruickshank, and Carley 2021). Characterizing, understanding, and eventually forecasting how online mediated interactions influence social cohesion and collective decision-making is part of an emerging field of study called social cybersecurity (Wang, Sun, and Zhu 2020; Carley 2020).

The detection of influence operations on social media is no easy task, as these inauthentic state-sponsored accounts put much effort into masking their activity, thus appearing indigenous to their target platform (Alizadeh et al. 2020). Frameworks for detection of coordinated activity based on high levels of common actions have been proposed and studied in the context of high profile events such as the 2016 US Presidential Elections and the 2021 ReOpen America protest campaigns (Weber and Neumann 2020; Magelinski, Ng, and Carley 2022). To handle the evolution of campaigns across time and the differences in influence operations by different groups, supervised machine learning frameworks leveraging content, networks, and user features help identify signals of campaigns (Varol et al. 2017). In order to associate possible sources to the detected influence campaigns, biases in word embedding vectors have been harnessed as a technique to detect target groups (Alizadeh et al. 2020).

Prior research used linguistic analyses to characterize the thematic content and evolution of part-of-speech and stylistic features during the 2016 US Presidential elections (Boyd et al. 2018) and detect signatures of state-sponsored influence operations through word embeddings (Toney et al. 2021).

Image Analysis. An image speaks a thousand words. Much of the work on image analysis in social media has been in meme identification and comparison (Beskow, Kumar, and Carley 2020), image sentiment and emotion analysis (Wang and Li 2015; Ng and Carley 2021), personality analysis through profile pictures (Liu et al. 2016), and profiling natural disasters and epidemics (Schreck and Keim 2013; Nguyen et al. 2017). Other image analysis work dwells on breaking down YouTube and TikTok videos into image frames and leveraging their comments to characterize the propagation of themes (Hussain et al. 2018; Ng et al. 2021).

Recent work analyzing themes in image clusters used to incite freedom fighters (Ng and Carley 2021) and characterizing the content shared by state-sponsored Russian trolls

(Zannettou et al. 2020) marks a shift in research to leverage image analysis to better understand influence operations. In particular, Zannettou et al. (2020) constructed network graphs representing the similarity of images of inauthentic Russian accounts and observed that these accounts are highly efficient in spreading political images in attempts to sow discord online during dividing events.

Image similarity assessment techniques typically involved distance-based functions such as the Euclidean distance between vectors containing information about invariant image features (Di Gesu and Starovoitov 1999). Image feature signatures have been constructed based on texture, shape, and color of objects (Kelly, Cannon, and Hush 1995). Using a combination of image feature values that include grayscale, color, motion, orientation, etc., Manovich (2012) compared one million digital comic book images and clustered them into themes. We build on this legacy and harness image similarity analysis techniques to characterize image coordination during influence operations. More advanced techniques include simulation to find image variants (Morel and Yu 2009). We build on this rich legacy and harness image similarity metrics to characterize image coordination during influence operations.

Methodology

Dataset. The images are obtained from the inauthentic accounts dataset from the `authenticity.training` data from PhoMemes 2022 Challenge¹. These images come from state-linked information operations originating from China, Iran, Russia, and Venezuela identified by Twitter Transparency Center (Twitter 2021). This dataset contains 3801 images from China, Iran, Russia, and Venezuela.

We do not remove repeated images because the presence of repeated images can indicate coordination tactics, i.e., amplification or diffusion of images. However, in the absence of additional metadata about the users or timestamp of the images, we thus focus on the image similarity coordination structure.

Ethics. We worked with the publicly available PhoMemes dataset and did not try to deanonymize account information.

Image Feature Extraction. We implement image feature extraction with Python Tensorflow. Each image is resized to a tensor with (244, 244, 3) dimensions as an input to the pre-trained ResNet50 image model (He et al. 2016). The image model thus represents the image in terms of a 2048-dimensional vector via 50-layers of convolutional neural networks trained on more than a million images with 1000 object categories.

Image-Image Network. We construct a k-Nearest Neighbors Image-Image similarity network. For each image, we perform a pairwise comparison between the other images in the comparison set by calculating the Euclidean distance between the two image vectors. To ensure flexibility of our method, we set the number of clusters k to be $k = \log N$, where N is the number of images of the comparison set, thus adjusting k to the size of the data. A set of k images is said to be similar based on shared image features.

¹<https://phomemes.github.io>

We then constructed an Image-Image network graph $G_I = (V_I, E_I)$ where V_I is the set of nodes $\{V_{I_1}, V_{I_2} \dots V_{I_n}\}$ representing images and $E_I(V_{I_j}, V_{I_k})$ represents an edge between two image nodes j and k if the images j and k are similar to each other. In short, a node represents an image, and edges represent the extent of similarity between two adjacent images. We lay G_I out with a ForceAtlas2 layout. We did this for images within each country and across all images in the dataset. We interpret networks in terms of (a) clustering coefficient, a measure of the degree to which the nodes in a graph tend to cluster together; and (b) fragmentation, the measure of the disconnectivity of the nodes in the graph.

Image Clustering. We perform network-based image clustering to identify the key narratives spread through the images. We first segregated the Image-Image network graph using the Louvain clustering algorithm (Blondel et al. 2008). Then we identify the top six clusters with the most number of images and manually interpret each image cluster.

Account-Account Network. To represent the coordination of accounts through similar images, we folded the Image-Image network to form an Account-Account network for each country and all the accounts in the dataset. In this, the Account-Account network graph $G_A = (V_A, E_A)$ where V_A is the set of nodes $\{V_{A_1}, V_{A_2} \dots V_{A_n}\}$ representing accounts and $E_A(V_{A_j}, V_{A_k})$ represents an edge between two account nodes j and k if they have images that are similar to each other. In short, a node represents an account, and edges represent the extent of similar images of two adjacent accounts. We lay G_A out with a ForceAtlas2 layout.

Results

Our method produces noticeable results in terms of the resulting structure of the image-to-image network, account-to-account network, and the top image clusters.

Image-Image Network. We list the network level metrics in Table 1. We use these metrics to reason about the level of coordination in the use of images for influence operations. China has the highest clustering coefficient, indicating a higher coordinated use of images than Iran and Venezuela. All networks but Russia exhibit some fragmentation that may represent noise or attempts not to appear coordinated by state actors. The lack of fragmentation in Russia’s network means that every image is similar to at least one other image in the Russian network.

We visualize the Image-Image networks by country in Figure 1. Differences in the clustering coefficient metric can be observed through the network structures: China, Iran, and Venezuela’s networks can be visually segmented. In contrast, the network of Russia displays a single centralized cluster. The presence of isolate nodes that line the outer circle of the networks shows the fragmentation of the Image-Image networks and the image themes.

The image model and similarity measure picks up on general image features such as faces, scenes, and words (memes) in images. This aids us in uncovering visual features used by one or many actors in the same country/event.

From the Image-Image network graph of all countries, we observe an interwoven graph, showcasing that images from these four countries generally build on threads from each other. However, some threads are unique to specific countries, represented by the stray network lines sticking out of the core network.

Account-Account Network. In Figure 2 we display the Account-Account networks for each country. We find that the networks are densely connected, consistent with the nature of the data, which are accounts identified for coordinated inauthentic behavior by Twitter. Interestingly, we observed varying levels of connectivity between accounts, as evidenced by the edge weights. Accounts with stronger connectivity indicate a larger number of similar image pairs between them. Nodes that are strongly connected may help identify the main coordinators of the influence campaigns. Unfortunately, we do not have the account details in this dataset and are thus unable to identify additional characteristics of the central coordinating accounts.

Country-Country Network. We further represent the Image-Image and Account-Account networks for all countries in Figure 3 to showcase the image coordination interaction between all four countries. We observe that Iran, Russia, and Venezuela have strong ties between images and accounts, i.e., many images within these three countries are similar to each other. They have many accounts connected through image similarity.

Table 1: Network level metrics for Image-Image networks

Country	Images	Clustering Coefficient	Fragmentation
China	598	0.140	0.108
Iran	890	0.110	0.107
Russia	1106	0.113	0.000
Venezuela	1207	0.112	0.149

Image Clustering. Figure 4 displays representative images of the top six network clusters obtained through the Louvain clustering algorithm for each country’s network. We find that the top clusters reveal distinct and coherent topics for each country. In general, the clusters reveal country-specific themes and influence operations. We also find image clusters of memes – images with word overlay – highlighting the use of memes by inauthentic accounts to disseminate information and ideology, probably in a coordinated fashion.

Discussion

We utilize past forensic evidence and studies on social media behavior, training, and funding estimates as a guide to interpreting our results. We highlight common tactics, account interaction, and communication styles for each country in terms of our image and network analysis results.

China has a high level of coordination and vertical integration from political party to agencies contributing to influence operations (Bradshaw and Howard 2017). There have been estimates that there are over millions of well-trained personnel contributing to influence operations, and the presence of activity-based reward systems (Bradshaw and Howard 2017; Han 2015; King, Pan, and Roberts 2017).




Figure 1: Image-Image Network. Nodes represent images. While the images are small to fit the screen, this figure emphasizes the network structure derived from image similarity metrics. Edge widths represent the similarity between two images. The salient clusters are depicted in Figure 4.

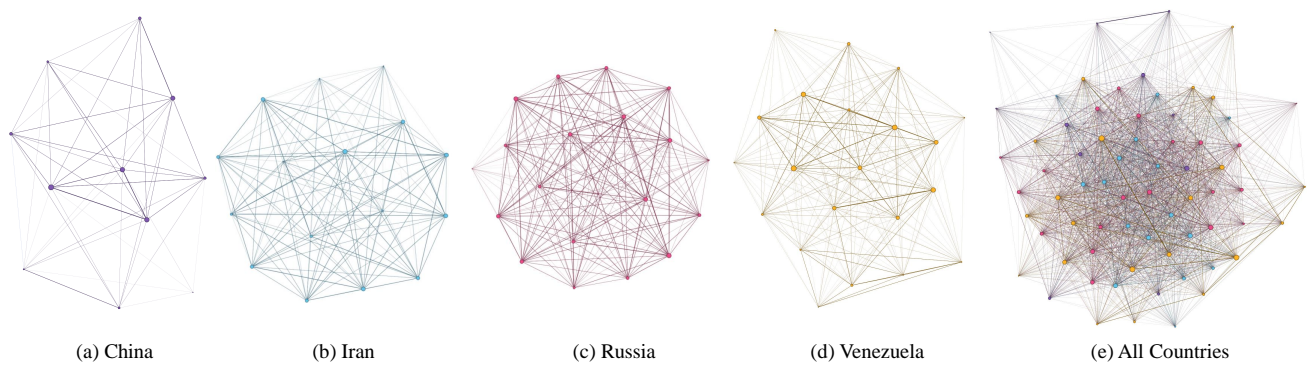


Figure 2: Account-Account Network. Nodes represent inauthentic accounts. Edge widths represent the number of similar images between the two accounts. Adjacent nodes with thicker edges represent a higher number of similar images shared between the two accounts.





Figure 3: This figure provides the country by country similar image network. In (a), we size nodes by the total number of images, and in (b), we size nodes by the number of unique authors. In (a) and (b), link weights are sized by the number of similar images shared between countries.

(1)



(2)



(3)



(4)



(5)



(6)



(a) China



(b) Iran

(1)



(2)



(3)



(4)




(5)



(6)



(c) Russia



(d) Venezuela

Figure 4: Representative images of the top six image clusters obtained through the Louvain clustering algorithm

Tactics of Chinese influence operations include pro-China messaging, suppressing or censoring messages that counter China's interests, and shifting attention away from essential issues (Bradshaw and Howard 2019; Singer and Brooking 2018).

Seen through this lens of knowledge about Chinese influence operations, we find three distinct lines of effort in the resulting image clusters for China. First, in row 2, we find memes mocking Guo Wengui, an exiled Chinese billionaire. Second, in row 6, images of the same Chinese billionaire are accompanied by phrases of repentance. Finally, in rows 1, 3, 4, and 5, we find images of lifestyle themes (luxury goods, food, travel, and pop stars). The use of lifestyle type images may be an attempt by influence campaign organizers to leverage popular topic groups of China's rising middle class in an attempt to build a following around its main influence operation (Barton, Chen, and Jin 2013), but more importantly, to brand its influence efforts as direct communication with the global public (Bund 2021). The Image-Image network's structure and high clustering coefficient may correlate to high-level coordination and integration of influence operations.

Iran exhibits a fraction of China's capacity for conducting influence operations, with approximately 20,000 practitioners (Bradshaw and Howard 2017). Tactics of Iranian influence operations include pro-government messaging, attacks on the opposition, and suppressing or censoring messages that counter Iran's interests (Bradshaw and Howard 2019). Image-based tactics include suppressing political dissidents using political hate speech, vulgar speech, counter-speech, and religion & societal topics (Kargar and Rauchfleisch 2019).

Russia has a well-funded, well-trained, and highly coordinated influence operation apparatus (Bradshaw and Howard 2017). On top of utilizing similar tactics as discussed with Chinese and Iranian capabilities, the Russian influence operations also focus on driving division and polarization in specific countries or alliances (Arif, Stewart, and Starbird 2018; Thornton 2015).

The Russian Image-Image network is highly connected, with no isolated images, which may further indicate the high-level coordination in the countries' influence operations. Its Account-Account Network is also highly connected and has more edges within the accounts than the other countries' networks. In the resulting image clusters for Russia, we find three distinct lines of efforts. First, evidence of attempts to drive division within the US political landscape in rows 3 and 6 through distorting images of Democratic party candidates and promoting divisive topics in the US media. Second, there are military-based and political-based images in rows 1 and 2, possibly associated with the NATO alliance and its opposition. Finally, we see a line of effort linking lifestyle themes of food and travel, similar to the observations of Chinese influence efforts.

Venezuela employs multiple bridges of people for influence operations in roughly 500-person teams and utilizes pro-government messaging, attacking the opposition, distraction, and suppression tactics (Bradshaw and Howard 2019). There is also evidence that Venezuela is starting to

copy Russian approaches to influence operations (Lucas and Nimmo 2015).

The Venezuela Image-Image network features the most isolates, suggesting a good proportion of images are disconnected from the central messaging theme. Many of Venezuela's isolated images are images of "Breaking News," with different news headlines, suggesting one of the Venezuelan strategies is to tweet about breaking news to disguise themselves as a news feed.

There are varying levels of connectedness between nodes in the Account-Account network, which may result from the organizational structure of the Venezuelan influence operation personnel.

In terms of image clusters, we see three clear lines of effort. First, in rows 1 and 2, there are political memes. These memes follow a basic template of an image with words written in large white Impact font. In contrast, memes by China have cartoon characters; those by Iran are comic-looking, suggesting that Venezuela's effort in meme generation may be more primitive than the other countries. Second, in row 3, are images of key areas of political tensions like the South China Sea and the United States. Lastly, in rows 4 and 6 are images of protests and rallies.

All Countries In Figure 4 we find that Iran, Russia, and Venezuela have images and memes of U.S. politicians within their top six clusters, while China does not. After folding the Image-Image networks and Account-Account networks, we find strong connections, in terms of shared images, between Iran, Russia, and Venezuela and weaker connections between all three countries and China. Figure 3 displays the folded networks. The high correlation of similar images between Iran, Russia, and Venezuela could indicate coordination between countries or possibly sharing tactics for influence operations. Further research would be needed to substantiate this observation.

Limitations and future work. The primary limitation of this study is the PhoMemes challenge data set itself. The images are studied in isolation from the texts and other meta-data features of the tweets and users, thus losing essential knowledge to understand the intended influence operations better. Additionally, the data may suffer from selection bias and a lack of transparency on how Twitter identified the accounts as part of influence operations.

The limitations in data provide excellent avenues for future work in studying the use of images in combination with other features in the Tweets. In particular, with tweet and user meta-data, we will construct directed network graphs to identify the diffusion of images across user-to-user all communication networks. Another branch of analysis will entail analyzing influence operations through image networks by characterizing observed changes across time.

Conclusion

In this study, we characterize the coordinated image dissemination strategy as part of the influence operations of four countries: China, Iran, Russia, and Venezuela. Using image similarity techniques, we construct Image-Image network graphs to identify the coordinated clusters of images.

We observe that images from Russia have a single centralized messaging, while images from other countries spout multiple messaging efforts, resulting in isolate nodes in the Image-Image network graphs. We interpret our network and image cluster observations through the lens of studies about the structure and organization of influence operations in each country, identifying distinct lines of effort in resulting image clusters for each country. By folding the Image-Image network graph, we construct Account-Account network graphs to represent the coordination between accounts through the use of similar images. We observe different levels of connectedness in the network graphs, possibly reflecting the funding and training of the influence operation apparatus of each country.

Grouping images and understanding their links are crucial in large-scale media analytics, aiding deep dives into data subsets for further analysis. Our work provides an unsupervised method to find and characterize groups of accounts and images quickly. This method is a critical first step in a more extensive analysis pipeline for investigating inauthentic coordination activity. We hope that our work provides a preliminary investigation method to aid researchers in understanding the strategic spread of visuals on social media as part of coordinated influence operations.

Acknowledgments

We would like to thank Jon Storrick from the CASOS, CMU lab for helping us with building a new feature to handle the image visualization on the network visualization software ORA. The research for this paper was supported in part by the Knight Foundation and the Office of Naval Research grant N000141812106 and by the center for Informed Democracy and Social-cybersecurity (IDeaS) and the center for Computational Analysis of Social and Organizational Systems (CASOS) at Carnegie Mellon University. The views and conclusions 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 US Government.

References

- Alizadeh, M.; Shapiro, J. N.; Buntain, C.; and Tucker, J. A. 2020. Content-based features predict social media influence operations. *Science advances*, 6(30): eabb5824.
- Arif, A.; Stewart, L. G.; and Starbird, K. 2018. Acting the part: Examining information operations within# BlackLivesMatter discourse. *Proceedings of the ACM on Human-Computer Interaction*, 2(CSCW): 1–27.
- Badawy, A.; Addawood, A.; Lerman, K.; and Ferrara, E. 2019. Characterizing the 2016 Russian IRA influence campaign. *Social Network Analysis and Mining*, 9(1): 1–11.
- Bakshy, E.; Hofman, J. M.; Mason, W. A.; and Watts, D. J. 2011. Everyone’s an influencer: quantifying influence on twitter. In *Proceedings of the fourth ACM international conference on Web search and data mining*, 65–74.
- Barton, D.; Chen, Y.; and Jin, A. 2013. Mapping China’s middle class. *McKinsey Quarterly*, 3(2013): 54–60.
- Bergh, A. 2020. Understanding Influence Operations in Social Media. *Journal of Information Warfare*, 19(4): 110–131.
- Beskow, D. M.; Kumar, S.; and Carley, K. M. 2020. The evolution of political memes: Detecting and characterizing internet memes with multi-modal deep learning. *Information Processing & Management*, 57(2): 102170.
- Blondel, V. D.; Guillaume, J.-L.; Lambiotte, R.; and Lefebvre, E. 2008. Fast unfolding of communities in large networks. *Journal of statistical mechanics: theory and experiment*, 2008(10): P10008.
- Boyd, R. L.; Spangher, A.; Fourney, A.; Nushi, B.; Ranade, G.; Pennebaker, J.; and Horvitz, E. 2018. Characterizing the internet research agency’s social media operations during the 2016 US presidential election using linguistic analyses.
- Bradshaw, S.; and Howard, P. 2017. Troops, trolls and troublemakers: A global inventory of organized social media manipulation.
- Bradshaw, S.; and Howard, P. N. 2019. The global disinformation order: 2019 global inventory of organised social media manipulation.
- Bund, J. 2021. Finding China’s Edge: Engineering Influence Operations within the Limits of Social Media Platform Rules. Technical report, ETH Zurich.
- Carley, K. M. 2020. Social cybersecurity: an emerging science. *Computational and mathematical organization theory*, 26(4): 365–381.
- Di Gesu, V.; and Starovoitov, V. 1999. Distance-based functions for image comparison. *Pattern Recognition Letters*, 20(2): 207–214.
- Han, R. 2015. Manufacturing consent in cyberspace: China’s “fifty-cent army”. *Journal of Current Chinese Affairs*, 44(2): 105–134.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Hussain, M. N.; Tokdemir, S.; Agarwal, N.; and Al-Khateeb, S. 2018. Analyzing disinformation and crowd manipulation tactics on YouTube. In *2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, 1092–1095. IEEE.
- Kargar, S.; and Rauchfleisch, A. 2019. State-aligned trolling in Iran and the double-edged affordances of Instagram. *New media & society*, 21(7): 1506–1527.
- Kelly, P. M.; Cannon, T. M.; and Hush, D. R. 1995. Query by image example: the comparison algorithm for navigating digital image databases (CANDID) approach. In *Storage and Retrieval for Image and Video Databases III*, volume 2420, 238–248. International Society for Optics and Photonics.
- King, G.; Pan, J.; and Roberts, M. E. 2017. How the Chinese government fabricates social media posts for strategic distraction, not engaged argument. *American political science review*, 111(3): 484–501.

- Larson, E. V.; Darilek, R. E.; Gibran, D.; Nichiporuk, B.; Richardson, A.; Schwartz, L. H.; and Thurston, C. Q. 2009. Foundations of effective influence operations: A framework for enhancing army capabilities. Technical report, RAND ARROYO CENTER SANTA MONICA CA.
- Liu, L.; Preotiuc-Pietro, D.; Samani, Z. R.; Moghaddam, M. E.; and Ungar, L. 2016. Analyzing personality through social media profile picture choice. In *Tenth international AAAI conference on web and social media*.
- Lucas, E.; and Nimmo, B. 2015. Information Warfare: What Is It and How to Win It, CEPA INFOWAR PAPER No. 1. Center for European Policy Analysis, Washington, DC.[online] <http://cepa.org/sites/default/files/Infowar%20Report.pdf>.
- Magelinski, T.; Ng, L.; and Carley, K. 2022. A Synchronized Action Framework for Detection of Coordination on Social Media. *Journal of Online Trust and Safety*, 1(2).
- Manovich, L. 2012. How to compare one million images? In *Understanding digital humanities*, 249–278. Springer.
- Morel, J.-M.; and Yu, G. 2009. ASIFT: A new framework for fully affine invariant image comparison. *SIAM journal on imaging sciences*, 2(2): 438–469.
- Ng, L. H.; and Taeihagh, A. 2021. How does fake news spread? Understanding pathways of disinformation spread through APIs. *Policy & Internet*, 13(4): 560–585.
- Ng, L. H. X.; and Carley, K. M. 2021. Bot-Based Emotion Behavior Differences in Images During Kashmir Black Day Event. In Thomson, R.; Hussain, M. N.; Dancy, C.; and Pyke, A., eds., *Social, Cultural, and Behavioral Modeling*, 184–194. Cham: Springer International Publishing. ISBN 978-3-030-80387-2.
- Ng, L. H. X.; Cruickshank, I.; and Carley, K. M. 2021. Coordinating Narratives and the Capitol Riots on Parler. *CoRR*, abs/2109.00945.
- Ng, L. H. X.; Tan, J. Y. H.; Tan, D. J. H.; and Lee, R. K.-W. 2021. Will You Dance to the Challenge? Predicting User Participation of TikTok Challenges. In *Proceedings of the 2021 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, ASONAM '21, 356–360. New York, NY, USA: Association for Computing Machinery. ISBN 9781450391283.
- Nguyen, D. T.; Ofli, F.; Imran, M.; and Mitra, P. 2017. Damage assessment from social media imagery data during disasters. In *Proceedings of the 2017 IEEE/ACM international conference on advances in social networks analysis and mining 2017*, 569–576.
- Rowett, G. 2018. The strategic need to understand online memes and modern information warfare theory. In *2018 IEEE International Conference on Big Data (Big Data)*, 4437–4442. IEEE.
- Schill, D. 2012. The Visual Image and the Political Image: A Review of Visual Communication Research in the Field of Political Communication. *Review of Communication*, 12(2): 118–142.
- Schreck, T.; and Keim, D. 2013. Visual Analysis of Social Media Data. *Computer*, 46(5): 68–75.
- Singer, P. W.; and Brooking, E. T. 2018. *LikeWar: The weaponization of social media*. Eamon Dolan Books.
- Stout, M. 2017. Covert action in the age of social media. *Georgetown Journal of International Affairs*, 94–103.
- Thornton, R. 2015. The changing nature of modern warfare: responding to Russian information warfare. *The RUSI Journal*, 160(4): 40–48.
- Toney, A.; Pandey, A.; Guo, W.; Broniatowski, D.; and Caliskan, A. 2021. Automatically Characterizing Targeted Information Operations Through Biases Present in Discourse on Twitter. In *2021 IEEE 15th International Conference on Semantic Computing (ICSC)*, 82–83. IEEE.
- Twitter. 2021. Information operations - Twitter Transparency Center.
- Varol, O.; Ferrara, E.; Menczer, F.; and Flammini, A. 2017. Early detection of promoted campaigns on social media. *EPJ data science*, 6: 1–19.
- Wang, Y.; and Li, B. 2015. Sentiment analysis for social media images. In *2015 IEEE international conference on data mining workshop (ICDMW)*, 1584–1591. IEEE.
- Wang, Y.; Tahmasbi, F.; Blackburn, J.; Bradlyn, B.; Cristofaro, E. D.; Magerman, D.; Zannettou, S.; and Stringhini, G. 2020. Understanding the Use of Fauxtography on Social Media. arXiv:2009.11792.
- Wang, Z.; Sun, L.; and Zhu, H. 2020. Defining social engineering in cybersecurity. *IEEE Access*, 8: 85094–85115.
- Weber, D.; and Neumann, F. 2020. Who's in the gang? Revealing coordinating communities in social media. In *2020 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, 89–93. IEEE.
- Zannettou, S.; Caulfield, T.; Bradlyn, B.; De Cristofaro, E.; Stringhini, G.; and Blackburn, J. 2020. Characterizing the use of images in state-sponsored information warfare operations by russian trolls on twitter. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 14, 774–785.