Conversations around Organizational Risk and Insider Threat

Luke Osterritter
CASOS, Insitute for Software Research
Carnegie Mellon University
Pittsburgh, PA, USA
losterritter@cmu.edu

Kathleen M. Carley
CASOS, Insitute for Software Research
Carnegie Mellon Univeristy
Pittsburgh, PA, USA
kathleen.carley@cs.cmu.edu

***Abstract*— Organizational risk and resilience as well as insider threat have been studied through the lenses of socio-psychological studies and information and computer sciences. As with all disciplines, it is an area in which practitioners, enthusiasts, and experts discuss the theory, issues, and solutions of the field in various online public forums. Such conversations, despite their public nature, can be difficult to understand and to study, even by those deeply involved in the communities themselves. Who are the key actors? How can we understand and characterize the culture around such communities, the problems they face, and the solutions favored by the experts in the field? Which narratives are being created and propagated, and by whom – and are these actors truly people, or are they autonomous agents, or "bots"?**

In this paper, we demonstrate the value in applying dynamic network analysis and social network analysis to gain situational awareness of the public conversation around insider threat, nation-state espionage, and industrial espionage. Characterizing public discourse around a topic can reveal individuals and organizations attempting to push or shape narratives in ways that might not be obvious to casual observation. Such techniques have been used to great effect in the study of elections, the COVID-19 pandemic, and the study of misinformation and disinformation, and we hope to show that their use in this area is a powerful way to build a foundation of understanding around the conversations in the online public forum, provide data and analysis for use in further research, and equip counter insider threat practitioners with new insights.

Keywords—insider threat, organizational risk, dynamic network analysis, social cybersecurity, online social networks

I. Introduction

Studying the risk posed by insiders can be difficult. An insider threat is most often defined as "a current or former employee, contractor, or business who has or had authorized access to an organization's network, system, or data and intentionally exceeded or misused that access in a manner that negatively affected the confidentiality, integrity, or availability of the organization's information or information systems," [1] a definition from the CERT Guide to Insider Threats. This definition works for threats across multiple types of organizations.

ASONAM '21, November 8–11, 2021, Virtual Event, Netherlands

ACM ISBN 978-1-4503-9128-3/21/11.
https://doi.org/10.1145/3487351.3492721

However, the way that one might discuss insider threat may have some difference and nuance depending on the type of threat, whether in the corporate world, or regarding nation states and espionage. While studying insider threat from the point of view of a specific organization or case study is valuable, there may also be value in studying the discourse surrounding threats in public online forum. Given this, we seek to answer whether dynamic network analysis can be used to discover the nature of public conversations around insider threat and related organizational threat.

We believe such exploration will be of value both to researchers and practitioners in the field. Insider threat researchers often survey groups of practitioners across organizations in order to elicit both expert feedback, and to characterize how insider threat programs are performing, what problems they face, and what solutions they employ [2],[3]. Analyzing social media conversation is one alternative, and in fact likely complimentary, way of garnering such information. It should be noted that the information collected will not be only from experts in the field, but by understanding how concepts are linked, what problems exist, and what solutions are most frequently discussed, it is possible to assess a "wisdom of the crowd" consensus, much in the way one might select the best solutions that are highly voted upon on social media sites like Reddit, StackExchange, or GitHub. In this sense as well, an understanding of what is being discussed is of value to training and education - what problems exist? How are people dealing with them? What are the best ways to deal with them?

The tools and techniques used to collect and analyze public forum conversations have potentially novel applications to open-source intelligence (OSINT) practitioners. Social network analysis is already recognized as one of the many available OSINT tools [4],[5], while dynamic network analysis techniques combine network analytics and semantical analytics, affording an understanding of both network structure and content. The use of such tools has been proven through repeated studies within just the last few years; these methods have been applied to terrorism and online extremism [6], hate speech and COVID-19 [7], fake news, misinformation, and disinformation [8] and recent elections [9]. Further, it may be useful to characterize public forum discussion before diving in to any one instance of targeted study, as the same techniques can be applied to study of specific groups or organizations – for instance, finding autonomous agents pushing narratives in the general

conversation could lead to further study of specific nation-state actors targeting commercial industry or other nation-state actors.

Finally, we believe the collection and synthesis of public forum data can help to inform and validate future research in a novel way, specifically with regard to applications in computational modeling. The results of studies like this can provide real-world data to inform models of insider risk and organizational resilience, while also providing a valuable opportunity for empirical validation. This is especially important for this topic, as it can be difficult to gather sufficient, publicly available data on instances of insider attacks or leakage. Having data on the types of attacks at different types of organizations perpetrated by different classes of individuals could then be used to inform, for instance, a system dynamics model of attack occurrence, or an agent-based model of interventions across different types of organizations.

II. RELATED WORK

As one of the largest available public forums, the social media site Twitter has been the subject of cyber-attacks resulting in the unauthorized access of prominent individuals and, subsequently, the focus of cyber security research and development [10]. As an extension of information available on the internet, social media joins traditional web sites as being a collection of potential resources for oppositions research and cyber-espionages [11]. Indeed, Twitter has potential value both in a more traditional cyber threat sense in the targeting of its users' accounts, but also a potential value for what is now being referred to as social espionage – specifically, the gathering of information on rivals via social media [12].

Social network analysis as a discipline has matured greatly, and appropriate techniques for its use in research have been codified in literature [13]. As a social network, Twitter data lends itself to study via social network analysis and statistical measures [14]. Such methods allow researchers to explore online conversations and map a social landscape, and to study a variety of issues endemic to a platform, such as polarization regarding political climates [15]. They allow for the discovery of users with malicious and possibly automated agents by analyzing the characteristics of known bad actors [16], and to further understand how bots and "Socialbot Networks" can be used to sway global conversations and public opinion for online extremist groups [17]. Social media discourse can also be leveraged to gain situational awareness of ongoing cyber-attacks [18]. Study of online discussions, and analysis of sentiment with regards to nation-states, can be a predictor of future cyber mediated attacks [19].

The study of Twitter data has increased year-by-year and occurs across many academic domains utilizing a wide variety of methodologies [20]. Per Zimmerman, "Studies from computer science, information science, and communications dominate, but a growing interest is evident from the domains of business, economics, education, medicine, political science, and sociology." Methodologies vary widely, but can include content analysis, sentiment analysis, and social network analysis [21]. Williams found also that the issues studied via Twitter in literature vary widely. These include influence via communication, emergency and disaster response, culture and politics, health and medical issues, and the use of automated tweeters or bots.

Importantly, gaining an understanding of a social network can afford the ability to strengthen or destabilizes said network – coupled with computational modeling, one is able to run what-if analyses on the removal of key agents in order to guide actions and inform policy surrounding social groups [22]. Kumar and Carley have gone further, using a pipeline approach wherein Twitter data is collected and augmented using machine learning tools [19]. Such tools, notably the Bot-hunter tool [23] employed in this study, allow for the identification of fully or partially automated social media accounts. Such knowledge can be valuable when analyzing public discourse, specifically for the discovery of entities subtly attempting to sway a conversation in their favor.

This work is part of an emerging and swiftly growing field of study referred to as social cybersecurity. Social cybersecurity is "focused on the science to characterize, understand, and forecast cyber-mediated changes in human behavior, social, cultural and political outcomes, and to build the cyber-infrastructure needed for society to persist in its essential character in a cyber-mediated information environment under changing conditions, actual or imminent social cyber-threats. [24]". This area of study affords tools, tactics, procedures, and policies for the analysis, prediction, and mitigation of social cyber-attacks [25]. The spread of information regarding insider threats via social media and the use of social media to conduct and defend against such adversarial actions is central to social cybersecurity. This paper, like other work in the area, employs high-dimensional network analytics to assess both connections among ideas and connections among actors.

III. DATA COLLECTION

As an initial step, we identified a set of twelve hashtags that would have some relation to a conversation around threats to an organization. When choosing terms, we believed it important to choose those that would return tweets that are strongly related to the topic of insider threat, but also wanted to consider terms that would allow us to retrieve conversations regarding threats in both a corporate and governmental context. As illustrated below, we have three categories of four hashtags each that allow us to collect a general conversation around insider threat as well as two specialized domains – corporate entities and nation-states.

These hashtags were chosen in an attempt to curate three separate categories within the insider threat field with as little overlap as possible, and the number of hashtags per grouping was kept equal to ensure symmetry for analytical purposes. The corpus is inclusive of many more relevant hashtags that could have been candidates for our initial collection, as some exploratory analysis was done early on to limit the amount of off-topic tweets returned from our queries. We believe this narrow scoping serves this work well, as it represents the first attempt to characterize the public conversations around this topic. That said, there has been work done on ontologies for insider threat [26],[27] that could inform future work in this area using different sets of hashtags into new groupings. We intend to explore this in future work.

TABLE I. TABLE 1. SET OF HASHTAGS USED FOR TWEET COLLECTION BY CONVERSATION CATEGORY

Category	Hashtags			
General	#insiderthreat	#insiderattack	#cyberespionage	#dataloss
Corporate	#industrialespionage	#tradesecrets	#embezzlement	#embezzling
Nation-state	#militarysecrets	#spy	#spying	#spies

To collect the relevant tweets, we used a Python package, twarc [28], which allows for using Twitter's search API to retrieve tweets based on a query, which are stored in a set JSON documents. We then imported these documents into the ORA-PRO software package [29],[30], which can parse Twitter-formatted social networks and affords powerful social network analysis and the ability to run statistical measures on social media data.

In the initial data collection, we retrieved 296 tweets over a one-week span. The nodesets included 142 unique agents, 208 hashtags, 302 total tweets, and 142 unique URLs. In the Agent x Agent all communication networks, there are 191 links between 141 agents, with a density of .0095.

After this initial collection, we collected for an additional seven weeks, for a total of eight non-contiguous weeks of data in each category, comprised of the months of March and July 2020. This collection includes a total of 46,535 tweets, 23,392 agents, 15,590 hashtags, and 8,446 URLs.

Once parsed into a meta-network within ORA-PRO, a total of 17 multi-modal networks become available for study: Agent x Agent networks of all communication, common hashtags, mentioned-by, quoted-by, reciprocal, replied-by and retweeted-by; an Agent x Hashtag network; Agent x Tweet (by sender); Tweet x Agent (mentions); Agent x URL; Hashtag x Hashtag co-occurrence; Tweet x hashtag; and Tweet x Tweet quoted by and replied-by.

We employed the Bot-hunter tool [23] to label twitter agents as engaging in bot-like, partially or fully automated activity. This tool was run on each batch of tweets collected, one file per hashtag, per week. For our purposes, a bot probability score of .75 or above resulted as an agent being labeled as having bot-like activity. While the agents in our dataset have been de-identified, those labeled as verified by Twitter or self-reported as news agencies have had their usernames retained in our reporting.

Bot-hunter was chosen as the tool to determine the presence of autonomous agents in the dataset for several reasons. Bot-hunter has several tiers of analysis, one of which is able to be run on data collected from the Twitter API and to perform an offline semantic assessment, rather than taking as input a list of usernames and querying the Twitter API itself. This has the benefit of not hitting the bottleneck of the Twitter API rate limits, can be run offline, and at a large scale on data of significant size. Autonomous agents on Twitter tend to go against the terms of service for the platform, and as such any tool that analyzes accounts directly from the API would be unable to assess terminated accounts, severely limiting our ability to assess whether such accounts were suspended because of bot-like behavior or other "Terms of Service" violations. Finally, bot-hunter is performant; recent benchmarking shows that bot-hunter tends to perform at or above the level of similar tools [31].

TABLE II. TABLE 2. PROPORTION OF ALL USERS TO THOSE DISPLAYING BOT-LIKE BEHAVIOR

Category	Agent Total	Bot-like behavior	Percentage bot-like
General	1927	288	14.95%
Corporate	2798	734	26.23%
Nation-state	19356	4330	22.37%
All Hashtags	23392	5136	21.96%

We report on each batched set of tweets by category as outlined in table 1, as well as on the entire data set, to explore the conversation both by the general topics of insider threat, corporate threats, and nation-state threats, and more broadly overall. In order to study each set uniformly we are use three main approaches for analysis: a network visualization of the hashtags that co-occur within the dataset; a "Super Spreaders" analysis, which identifies agents whose tweets are often retweeted, meaning their information is more often spread across the network; and a "Super Friends" analysis, which identifies agents exhibit reciprocal communications with other agents, meaning that they are active in two-way conversations throughout the network.

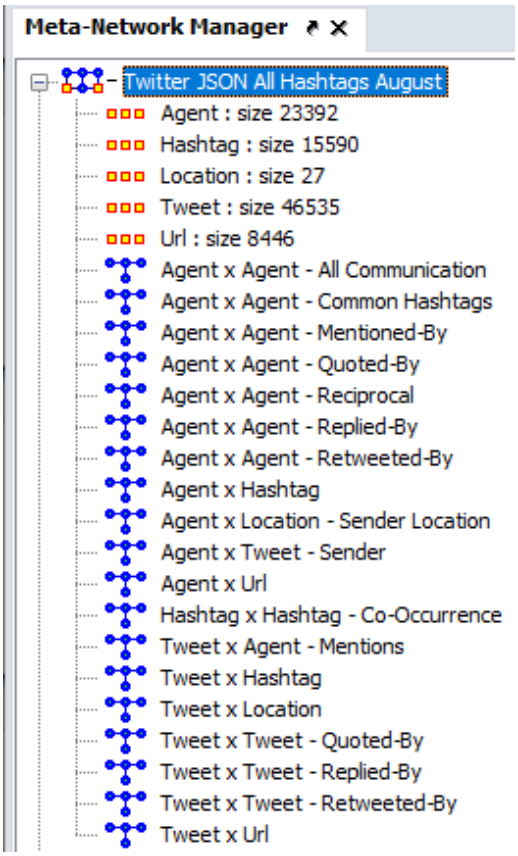

Fig. 1. Twitter nodesets and networks as parsed by ORA-PRO

We initially explored the hashtag co-occurrence network to gain a sense of which hashtags were common amongst the tweets that used the #insiderthreat tag. The Hashtag x Hashtag co-occurrence network is pictured below as an example of the initial exploratory data analysis done – labeled are the top 25 hashtags by usage count, and the colors represent Louvain grouping. This gives a birds-eye view of the conversation showing a color-coded representation of communities of hashtags; it is a visual representation of the nature of conversations that occurred in the data set.

This network reveals a conversation around insider threat, with sub-topics surrounding autonomous vehicles, cyber-crime, privacy, aviation concerns, healthcare, penetration testing techniques, intelligence, and the emerging COVID-19 pandemic. Studying the network in this manner is an effective way to gain situational awareness regarding public discourse and can inform the direction of further study.

IV. Dynamic Network Anlysis Results

A. Autonomous Agent Analysis

Given prior work on the preponderance of bots in public conversations, it would be naïve to think that the conversations collected in this study are being performed solely by human agents. This presents a problem regarding our situational awareness of the conversation; we would be unable to determine how organic the conversations within the data are without performing difficult and time-consuming manual labeling of each tweet. To this end, we have employed Bot-hunter [23] which leverages machine learning to assess which and how many agents in our data are acting in a bot-like manner. Coupled with the Twitter analysis features in ORA-PRO, we can than corroborate which actors important to the network are likely bots, and then further assess what their goals and motivations may be based on the content of their discourse.

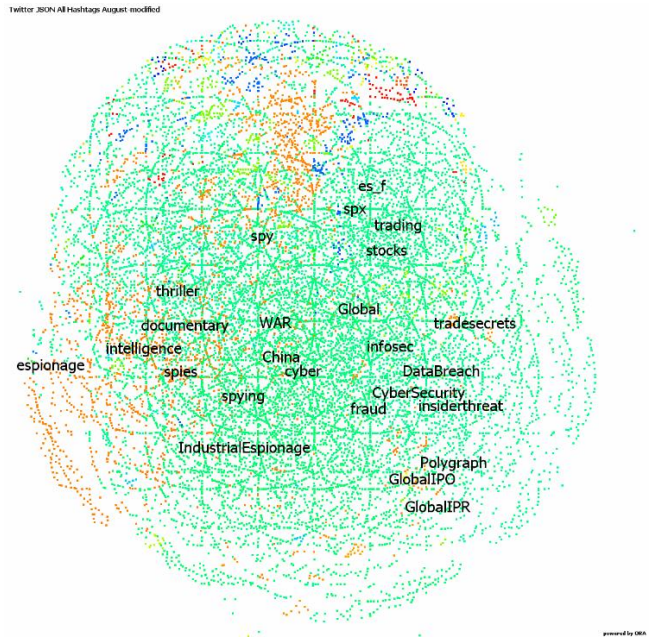

Fig. 2. Initial Hashtag x Hashtag co-occurrence network visualization for all data, colored by Louvain grouping

Fig. 3 is a graph of the Agent x Agent – All Communications network of the general insider threat hashtag grouping. A select number of verified agents have been labeled, shaded blue, and have diamond shapes. All nodes that are red in color and pentagonal in shape represent nodes that exhibited bot-like activity within the dataset. The remaining, teal-colored circles are otherwise normal users – neither verified by Twitter, nor labeled as bots by bot-hunter.

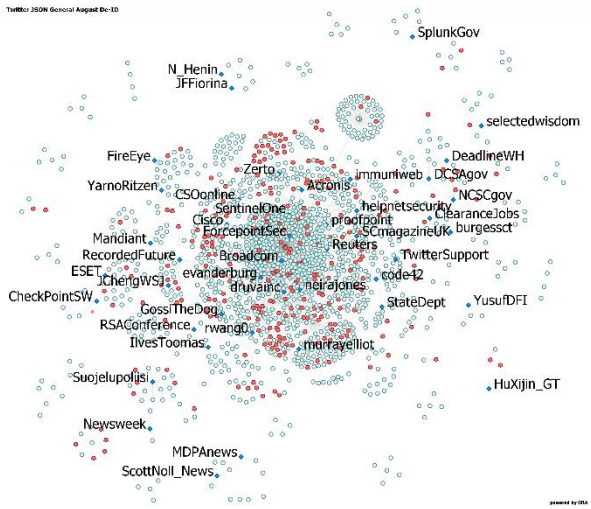

Fig. 3. Hashtag x Hashtag co-occurrence network across all tweets

B. Overview of Hashtags Across All Groups

The initial analysis was performed on all tweets collected for this study across all three groups and all 12 hashtags. Studying the data in aggregate gives a broad overview of how the whole conversation looks, reveals the nature of conversations occurring within the dataset, and shows who in the dataset is taking part in two-way conversations and spreading information across the network.

The network image (Fig. 4) shows that the overall conversation has a major hub in the #spy hashtag. Links with weights less than three standard deviations were removed to maintain readability of the graph – we can compare this to figure 2 to illustrate how it is possible to create a more granular view after the initial "bird's eye" view. Much of the conversation surrounding this hashtag refers to books, movies, video games, and other media. "Insider threat" connects with the concepts of cybersecurity and technology, and is reasonably close to conversations around industrial espionage, security, and privacy.

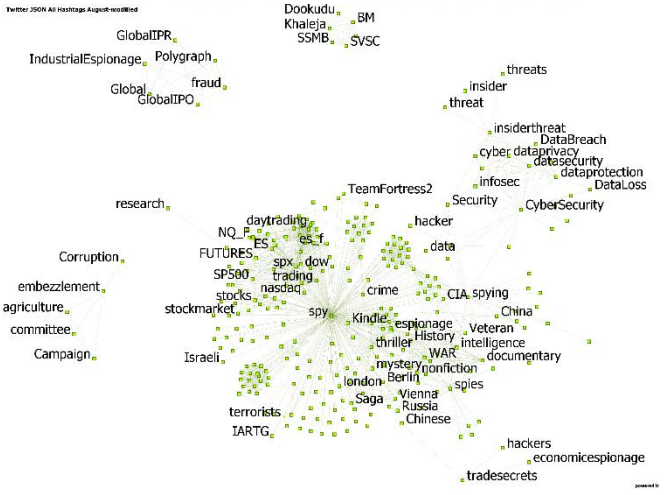

Fig. 4. Hashtag x hashtag co-occurrence network across all tweets

The analysis of Super Spreaders reveals accounts that largely exist to aggregate news articles, including three verified accounts. Verified accounts go through a verification process put forth by Twitter to confirm their identities, and often consist of news outlets, corporate accounts, and celebrities. One account was suspended, another has since been set to protected, and another displays some automated activity. Others are related to the marketing of works of fiction, and two others are related to a commercial software package. One such account identified as a bot was likely labeled as such due to repeated promotion of a novel release.

The Super Friends analysis is comprised largely of accounts having to do with comedy streaming media promotion, commercial software, economic markets, and fiction book promotion. Of the top ten super spreaders labeled in figure 5, all of these were found to be engaged in bot-like activity.

While these top agents are not closely related to the topic of interest for our study, the methodology of understanding the conversation using network analysis works well and studying each set of tweets individually will allow for deeper dive into topics of interest.

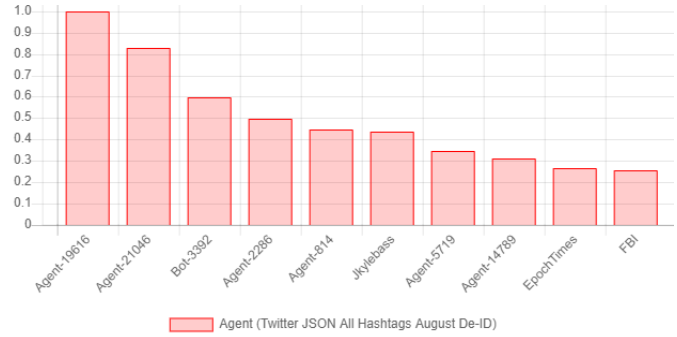

Fig. 5. Overall Super Spreaders

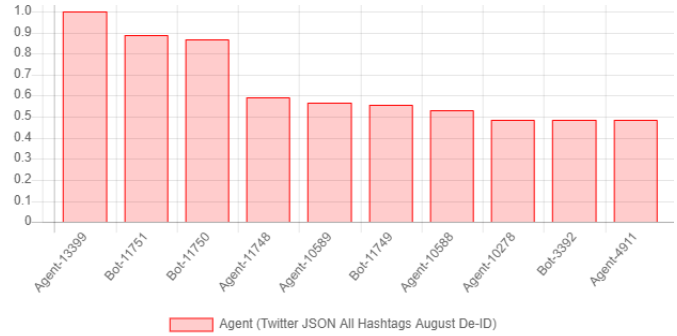

Fig. 6. Overall Super Friends

C. Insider Threat-based Hashtags

The network map of hashtag co-occurrence for this subgroup shows the conversations around the main insider threat hashtags. In fig. 7, links with weights less than one standard deviation were removed, as were isolate nodes. Cyber security, cyber risk, data loss, and InfoSec are all related concepts with heavy use in conversations. Conversations around AI, military, data protection, COVID19, specific security and technology companies, along with discussions surrounding specific advance persistent threats (APTs) are all present as well.

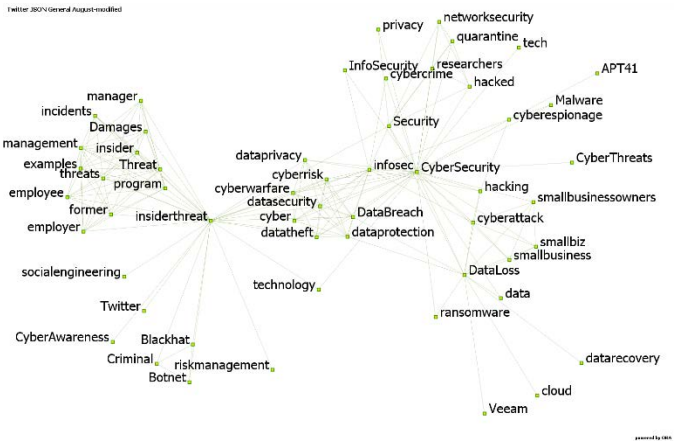

Fig. 7. Hashtag co-occurrence network, insider threat general hashtags

An analysis of super spreaders in the #insiderthreat tweet collection revealed the presence of one agent with a score much higher than any other agent. Further investigation reveals that this user was disabled sometime after the data was collected, indicating that they performed some actions that violated Twitter's terms of service.

Upon further analysis, this user appears to be a cyber-security professional – with a current web presence and an account on another social media platform for professionals, LinkedIn. A survey of this user's tweets as they appeared in the data collection, as well as through Google's search cache, do not reveal any specifically malicious messaging or behavior, though it does appear as though they were using an automated syndication program to send tweets automatically. Two possibilities are likely: either this account was disabled by Twitter for automatic activity, or the account was hijacked at some point, causing the suspension of the account and deletion of the offending tweets. It should be noted that the Bot-hunter score for the user in question here was well below the threshold, and that the algorithms used by Twitter to detect bot activity are not publicly available; the criteria used for bot detection and removal on the platform are not available to the public.

The other agents in these two reports consist of enterprise security companies, researchers, and news aggregators. Two of these accounts are verified, including two companies that are well-known in the insider threat space, FireEye and code42. This set of tweets aligns closely with the topic of insider and organizational threat, likely due to the lower levels of topic ambiguity with this grouping of collected hashtags.

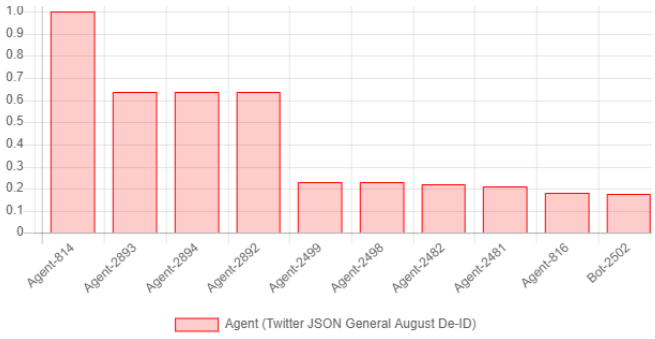

Fig. 8. Super Spreaders, insider threat general hashtags

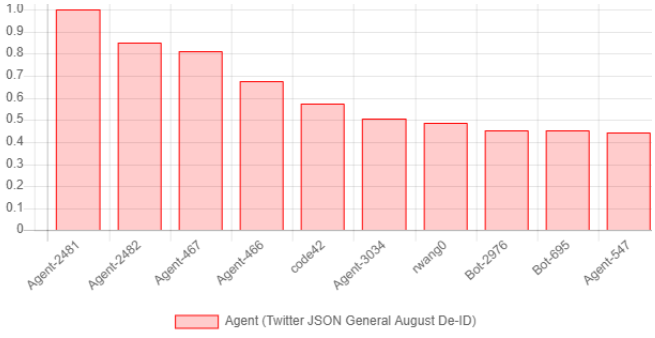

Fig. 9. Super Friends, insider threat general hashtags

D. *Corporate-based Hashtags*

In this set of hashtags, the network map shows that the concepts of industrial espionage, embezzlement, crime, and trade secrets feature prominently. Link weights greater than twice the mean value have been removed to increase the readability of the network visualization, as were all isolate nodes – those not having links to any other node in the network. The conversation around both trade secrets and embezzlement are varied, with the former touching on copyright, intellectual property, and AI, while the latter includes theft, regulation, and American politics. Indeed, there is at least some conversation not strongly linked to others that refers to a brand of American right-wing political sentiments.

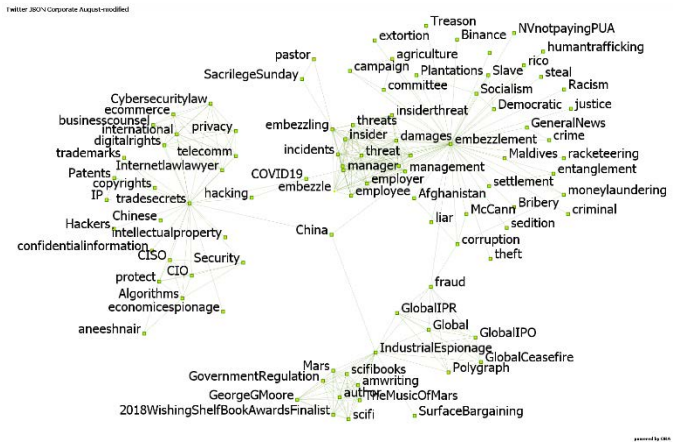

Fig. 10. Hashtag co-occurrence network, corporate hashtags

The top scoring account in the Super Spreader report is a verified account that espouses the freedom of movement and uniform application of EU laws for EU citizens. Interestingly, the runner up account uses Greek language and tweets at this first EU account with regularity, using anti-Europe rhetoric and the same set of hashtags. This account is still active and may be engaged in automated activity. Some of its tweets are in English while most are in Greek. Given its rhetoric, it is difficult to ascertain what its views or goals are beyond that it is against the EU and its member countries. Further analysis of this account could reveal a broader campaign against the Europeans governments.

Other accounts in these reports includes activists, news blogs and aggregators, law firms, financial and security consultants, political enthusiasts along the spectrum, and non-governmental organizations. Interestingly, an account that shows up in the super spreaders report and scored highly for bot-like behavior appears to be an account aligned with the American left posting rhetoric against the then-current incumbency. This part of the dataset has a bit more noise but is still very close to the main insider and organizational threat conversation.

E. *Nation-state-based Hashtags*

The network surrounding the nation-state-based hashtags reveals a conversation that is much broader than the other two individual categories. The Hashtag x Hashtag co-occurrence network in figure 13 has links weighted less than five standard deviations removed, as well as isolate nodes. The #spy hashtag

has conversations around much more than simply real-world spy activities. There is a large portion of the conversation surround economics and stock trading, and a great deal regarding books, movies, TV, and streaming entertainment. There are some concepts related specifically to both American and Canadian politics, as well as more generally NATO, Ireland, China, and Russia. Also present are topics related to COVID-19, 5G, and platforms like WhatsApp, Facebook, and TikTok.

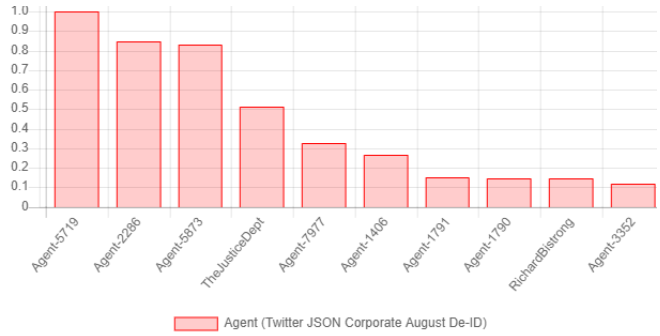

Fig. 11. Super Spreaders, corporate hashtags

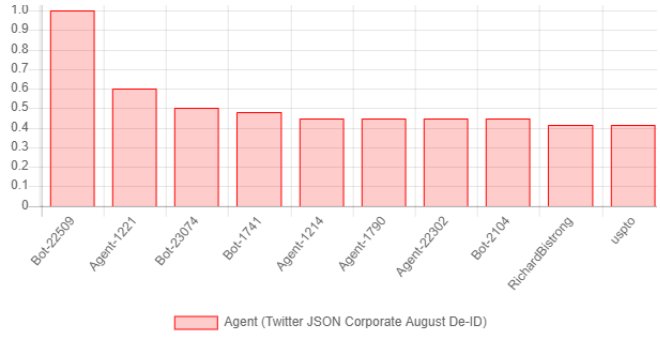

Fig. 12. Super Friends, corporate hashtags

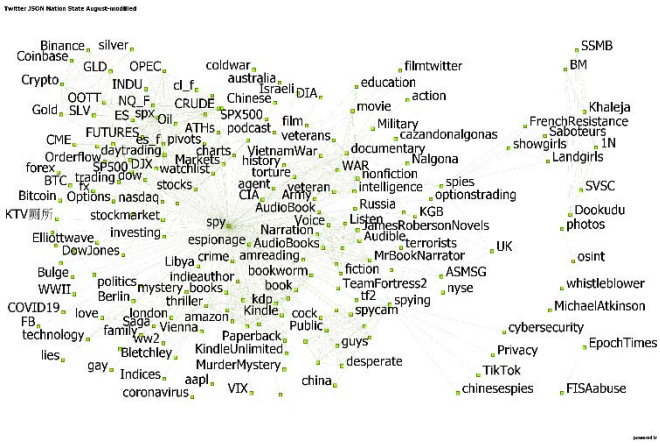

Fig. 13. Hashtag co-occurrence network, nation-state hashtags

The accounts ranked highly in the Super Spreaders and Super Friends reports include several accounts also present in the overall analysis. There are four total verified accounts. The types of accounts include movie promotion, news agencies, fiction authors, former intelligence workers and military, IT workers, software engineers, bloggers, commercial software, actors, economic enthusiasts, and web series promotion.

This category has the highest number of topics overall and is least representative of its intended topic retrieval – the conversation is there, but there is much noise to cut through. The top actors are mostly related to other non-national threat topics yet are very active and spreading large amounts of information, which affect the analysis of this group and the overall analysis. The bot activity in this dataset is largely related to the promotion of media, displaying heavy and repeated use of the same hashtags and rhetoric. This is similar to the bot activity seen in the collection of all hashtags, as the agents in this portion of the data proved to be strong in the data overall.

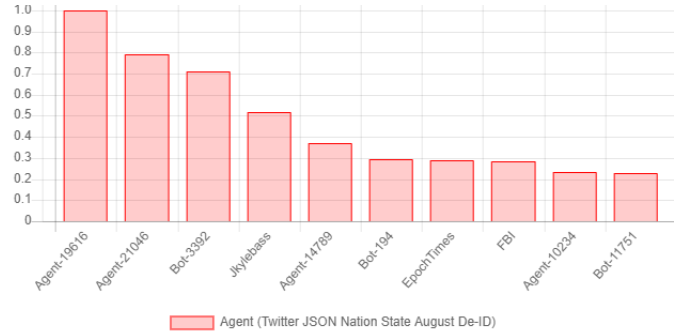

Fig. 14. Super Spreaders, nation-state hashtags

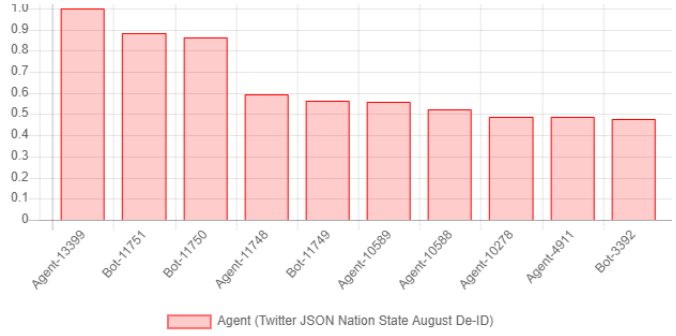

Fig. 15. Super Friends, nation-state hashtags

V. Discussion and Future Work

Collecting tweets surrounding hashtags around insider and organizational threat and analyzing this data by category allowed for gaining a clearer insight into how the public discussion around these topics exists over a period of time. Getting at the key actors related to the topic of study became more difficult as the collected hashtags became broader, though it is still of use to be able to see how the conversation takes shape in these related conversations. The nation-state data contained the discussion of interest, but also many other discussions that became harder to cut through – however, the categorization of hashtags by topic made this task much more successful. It may be useful to construct a corpus of Tweet based on existing ontological categorizations in literature relating to insider threat as discussed earlier in this paper – if such a corpus were of sufficient quality and its discourse sufficiently on-topic, it could be of greater value to researchers and practitioners familiar with such categorizations.

Overall, studying public discourse in this way resulted in the identification of key actors spreading information and engaging in conversation, understanding who those key actors are, and being able to find accounts that engaged in activity that was against the platform's terms of service. The ability to gain this type of situational awareness in cyber conversations is of high value to those interested in organizational research, cyber security, social cybersecurity, and policy.

This work would be supported by using natural language processing and further text analysis of the tweets in this dataset could allow for a deeper understanding of the public conversation, as extracting networks from those texts could be a complimentary way of analyzing the topics present and how they are connected.

As the data collected here is temporal in nature, it would be possible to analyze it as part of a dynamic meta-network. This would allow for looking at network structure and measures over time, as opposed to in aggregate as presented in this study. This dataset contains tweets from just before and just after the events of the COVID-19 pandemic began to be realized in America – studying it from this angle with temporal elements would allow for study of the global discourse before and during such events.

The hashtags present alongside those collected on are an opportunity for addition-al study and could be implemented into the data collection strategy for a future study. For instance, many advance persistent threats are present as hashtags in this data and studying the discussion around those could lead to a study of the technical tactics, tools, and procedures surrounding the carrying out of malicious insider at-tacks.

Finally, the use of other tools around identified key actors could allow for attribution of actors across platforms – that is to say, using the tools and techniques presented here to find actors engaged in the augmentation of public discourse, then using tools such as Maltego to identify these same actors on other platforms and observe their behavior in those forms could be a strong way of identifying and confirming attribution and intent.

VI. CONCLUSION

The network analysis techniques applied in this paper allow for gaining situational awareness of public discourse on online forums. We specifically looked at the discussion on the Twitter platform relating to insider threats broadly, as well as on the individual areas of insider threat in a general sense, corporate organization threats, and organizational threats to nation-states.

We employed machine learning tools to evaluate the presence of automated agents, or bots, participating in public discourse. This approach was helpful in cutting through the noise present in these discussions and understanding the state of the discussion in a more granular detail. These methods allowed for an understanding of this conversation, affording greater insight into the research problem, and informing future work in the area by research and practitioners. Additionally, these methods can be used on data surrounding other conversation to similar effect regarding data collection and situational awareness of the conversations contained therein.

ACKNOWLEDGMENT

This work was supported in part by the Office of Naval Research (ONR) Award N000141812106 and the Center for Computational Analysis of Social and Organization 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.

REFERENCES

[1] D. Cappelli, A. Moore, and R. Trzeciak, *The CERT Guide to Insider threats*. 2012.

[2] F. Greitzer, J. Purl, Y. M. Leong, and D. E. S. Becker, "SOFIT: Sociotechnical and Organizational Factors for Insider Threat," in *2018 IEEE Security and Privacy Workshops (SPW)*, May 2018, pp. 197–206. doi: 10.1109/SPW.2018.00035.

[3] A. P. Moore, T. M. Cassidy, M. C. Theis, D. Bauer, D. M. Rousseau, and S. B. Moore, "Balancing organizational incentives to counter insider threat," in *Proceedings - 2018 IEEE Symposium on Security and Privacy Workshops, SPW 2018*, May 2018, pp. 237–246. doi: 10.1109/SPW.2018.00039.

[4] P. L. Brantingham, "Computational Criminology," in *2011 European Intelligence and Security Informatics Conference*, 2011, pp. 3–3. doi: 10.1109/EISIC.2011.79.

[5] F. Tabatabaei and D. Wells, "OSINT in the Context of Cyber-Security," in *Open Source Intelligence Investigation: From Strategy to Implementation*, B. Akhgar, P. S. Bayerl, and F. Sampson, Eds. Cham: Springer International Publishing, 2016, pp. 213–231. doi: 10.1007/978-3-319-47671-1_14.

[6] M. C. Benigni, K. Joseph, and K. M. Carley, "Online extremism and the communities that sustain it: Detecting the ISIS supporting community on Twitter," *PLOS ONE*, vol. 12, no. 12, p. e0181405, Dec. 2017, doi: 10.1371/journal.pone.0181405.

[7] J. Uyheng and K. M. Carley, "Characterizing network dynamics of online hate communities around the COVID-19 pandemic," *Appl. Netw. Sci.*, vol. 6, no. 1, p. 20, Mar. 2021, doi: 10.1007/s41109-021-00362-x.

[8] M. Babcock, R. Villa-Cox, and K. M. Carley, "Pretending Positive, Pushing False: Comparing Captain Marvel Misinformation Campaigns," in *Disinformation, Misinformation, and Fake News in Social Media: Emerging Research Challenges and Opportunities*, K. Shu, S. Wang, D. Lee, and H. Liu, Eds. Cham: Springer International Publishing, 2020, pp. 83–94. doi: 10.1007/978-3-030-42699-6_5.

[9] T. Magelinski, M. Bartulovic, and K. M. Carley, "Canadian Federal Election and Hashtags that Do Not Belong," in *Social, Cultural, and Behavioral Modeling*, Cham, 2020, pp. 161–170.

[10] S. J. Stolfo, M. B. Salem, and A. D. Keromytis, "Fog Computing: Mitigating Insider Data Theft Attacks in the Cloud," in *2012 IEEE Symposium on Security and Privacy Workshops*, May 2012, pp. 125–128. doi: 10.1109/SPW.2012.19.

[11] M. S. Bressler and L. Bressler, "PROTECTING YOUR COMPANY'S INTELLECTUAL PROPERTY ASSETS FROM CYBER-ESPIONAGE," vol. 18, no. 1, p. 15, 2015.

[12] J. Salminen and W. Y. Degbey, "Social Media Espionage — A Strategic Grid," in *New Technology-Based Firms in the New Millennium*, vol. 11, Emerald Group Publishing Limited, 2015, pp. 261–274. doi: 10.1108/S1876-022820150000011020.

[13] S. Wasserman and K. Faust, "Social network analysis: Methods and applications," *Camb. Univ. Press*, 1994, doi: 10.1525/ae.1997.24.1.219.

[14] F. Morstatter, J. Pfeffer, H. Liu, and K. M. Carley, "Is the Sample Good Enough? Comparing Data from Twitter's Streaming API with Twitter's Firehose," pp. 400–408, 2013, doi: 10.1007/978-3-319-05579-4_10.

[15] M. a Smith, L. Rainie, I. Himelboim, and B. Shneiderman, "Mapping Twitter Topic Networks: From Polarized Crowds to Community Clusters," *Pew Res. Cent.*, no. February 20, pp. 1–57, 2014.

[16] W. Wei, K. Joseph, H. Liu, and K. M. Carley, "Exploring characteristics of suspended users and network stability on Twitter," *Soc. Netw. Anal. Min.*, vol. 6, no. 1, 2016, doi: 10.1007/s13278-016-0358-5.

[17] M. C. Benigni, K. Joseph, and K. M. Carley, "Bot-ivistm: Assessing Information Manipulation in Social Media Using Network Analytics," in *Emerging Research Challenges and Opportunities in Computational Social Network Analysis and Mining*, N. Agarwal, N. Dokoohaki, and S. Tokdemir, Eds. Cham: Springer International Publishing, 2019, pp. 19–42. doi: 10.1007/978-3-319-94105-9_2.

[18] R. P. Khandpur, T. Ji, S. Jan, G. Wang, C.-T. Lu, and N. Ramakrishnan, "Crowdsourcing Cybersecurity: Cyber Attack Detection using Social Media," in *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, Singapore Singapore, Nov. 2017, pp. 1049–1057. doi: 10.1145/3132847.3132866.

[19] S. Kumar and K. M. Carley, "Understanding DDoS cyber-attacks using social media analytics," in *2016 IEEE Conference on Intelligence and Security Informatics (ISI)*, Sep. 2016, pp. 231–236. doi: 10.1109/ISI.2016.7745480.

[20] M. Zimmer and N. J. Proferes, "A topology of Twitter research: disciplines, methods, and ethics," *Aslib J. Inf. Manag.*, vol. 66, no. 3, pp. 250–261, Jan. 2014, doi: 10.1108/AJIM-09-2013-0083.

[21] S. A. Williams, M. M. Terras, and C. Warwick, "What do people study when they study Twitter? Classifying Twitter related academic papers," *J. Doc.*, vol. 69, no. 3, pp. 384–410, Jan. 2013, doi: 10.1108/JD-03-2012-0027.

[22] K. M. Carley, J.-S. Lee, and D. Krackhardt, "Destabilizing networks," *Connections*, vol. 24, no. 3, pp. 79–92, 2002.

[23] D. M. Beskow and K. M. Carley, "Bot-hunter: A Tiered Approach to Detecting & Characterizing Automated Activity on Twitter," p. 8, Jul. 2018.

[24] K. M. Carley, G. Cervone, N. Agarwal, and H. Liu, "Social Cyber-Security," in *Social, Cultural, and Behavioral Modeling*, vol. 10899, R. Thomson, C. Dancy, A. Hyder, and H. Bisgin, Eds. Cham: Springer International Publishing, 2018, pp. 389–394. doi: 10.1007/978-3-319-93372-6_42.

[25] E. National Academies of Sciences, *A Decadal Survey of the Social and Behavioral Sciences: A Research Agenda for Advancing Intelligence Analysis*. 2019. doi: 10.17226/25335.

[26] D. L. Costa, M. J. Albrethsen, M. L. Collins, S. J. Perl, G. J. Silowash, and D. L. Spooner, "An Insider Threat Indicator Ontology," p. 87.

[27] F. L. Greitzer *et al.*, "Developing an Ontology for Individual and Organizational Sociotechnical Indicators of Insider Threat Risk," p. 9.

[28] "DocNow/twarc: A command line tool (and Python library) for archiving Twitter JSON." https://github.com/DocNow/twarc (accessed Apr. 13, 2020).

[29] N. Altman, K. M. Carley, and Raminga, "ORA User's Guide 2020," Carnegie Mellon University, School of Computer Science, Institute for Software Research, Pittsburgh, PA, Technical Report CMU-ISR-20-110, 2020.

[30] K. M. Carley, "ORA: A Toolkit for Dynamic Network Analysis and Visualization.," *Reda Alhajj Jon Rokne Eds Encycl. Soc. Netw. Anal. Min. Springer*, 2017.

[31] D. M. Beskow and K. M. Carley, "You Are Known by Your Friends: Leveraging Network Metrics for Bot Detection in Twitter," in *Open Source Intelligence and Cyber Crime: Social Media Analytics*, M. A. Tayebi, U. Glässer, and D. B. Skillicorn, Eds. Cham: Springer International Publishing, 2020, pp. 53–88. doi: 10.1007/978-3-030-41251-7_3.