Detecting Botnet Signals within Social Media Data

"Falsehood flies, and truth comes limping after it."
Jonathan Swift from The Examiner, Issue No. 14, Sept 11th 1710

Abstract

Detecting and elucidating botnets is an active area of research. The current state of botnet detection identifies automated features such as identical content, identical targets, coordination of message dispersal, and similar capabilities. Choreographed cross-platform botnets inject information warfare attack vectors into public discourse and are capable of dangerous destabilization. Sentiment models and bot detection methods alone are insufficient to detect and defend against weaponized artificial intelligence botnets that choreograph, amplify, and normalize messages of hate, anger, and violence. Process technologies are presented which detected botnets within terrorist Twitter data. Bot logic uses information camouflage in order to disguise intentions similar to World War II Nazi propagandists and Soviet-era practitioners of information warfare enhaced with reflexive control. A future effort is presented which strings together best of breed techniques into a composite classification algorithm in order to continually improve the discovery of malicious accounts and understand choreographed weaponized botnet dynamics which span social media platforms as well as recursive adversarial modeling.

Introduction

An information warfare arms race is underway, and the United States requires near real time capabilities to identify malicious social media accounts and control the information environment. In the early 21st Century, "peacetime is the decisive phase of operations (Cyber Endeavour Conference with non-attributional Chatham House Rules 2019).” Information Warfare (IW) domain scholars observe that technologies exist for scaling IW attack manually; the next stage of IW technology development is to automate it (Paul and Matthews 2016; Waltzman 2017). State and non-state actors are weaponizing artificial intelligence (AI) against the United States and Allies (Bicknell and Krebs 2019a; 2019b). Inexpensive social media platforms enable unprecedented IW campaigns which have turned the entire free world into a combat zone and every device into an attack vector delivery vehicle.

Using techniques analogous to corporate multi-channel marketing, enemies latch on to moments of crisis and exploit fault lines in our public discourse with highly effective IW campaigns by coupling ever sophisticated analyses along with choreographed multi-channel weaponized botnet social campaigns. This style of IW incorporates reflexive control (RC) theory which causes the controlled (attacked) party to make probabilistically pre-determined reflexive decisions which favor the attacker (Chotikul 1986; Thomas 2004; Bicknell and Krebs 2019a;

2019b), Incorporating cybernetics and game theory, RC IW attacks infect societies and foment destabilizing trajectories. Campaign effects are maximized by delivering curated information choreographed across multiple touchpoints thus exploiting fault lines in public discourse. Weaponized botnets, for example, cause controlled parties to be confused, mis-direct aggression at friendly forces, or exhibit mob behavior (Iyengar 2018; McLaughlin 2018). Bots are automated social media profiles designed to look like human users; they can be programed to complete functions such as tweeting, retweeting, liking, direct messaging or following/unfollowing other accounts. Bots may be fully automated or partially automated cyborg accounts run by a human but also augmented with automation to post faster, more frequently and more voluminously (Singer and Brooking 2018).

Processes underlie all complex naturally occurring phenomena (Whitehead 1979) and are an effective means to analyze time-domain relationships. Poorly understood or opaque processes are a national security intelligence gap. If processes of interest were visible and explicit, opportunities become available to strengthen, exploit or monitor these processes. Such processes include corporate, government, societal, transnational criminal, geography spanning, and others. State and non-state promulgators of social IW attacks are actively deploying countermeasures to avoid detection; yet, they are bound by natural laws and may be modeled with advanced process technologies. Understanding and identifying choreographed information warfare botnet activity is an active area of research and a pressing national security requirement ("DoD Early Detection of Information Campaigns by Adversarial State and Non-State Actors" 2019). Since known adversaries already think of their IW operations in terms of process modeling (Novikov and Chkhartishvili 2014), process mining AI is a framework especially well-suited to analyzing their activities and modeling their behavior.

More Cowbell Unlimited is working with Texas A&M's Cyber Security Center. Pioneering and computationally efficient process analyses reveal botnet signals latent within Twitter social media data. The results presented in this paper are a promising first step which will be expanded into a comprehensive cross-platform social media weaponized botnet ecosystem discovery and elucidation technique. Process mining is a powerful, nascent AI technique especially well-suited for analyzing data with an important temporal component, such as socially charged IW campaigns. More Cowbell Unlimited's cloud SaaS process mining software mines data and surface machine readable models of decision-making processes from various input formats.

Methodology

The goal of process mining is to turn event data into insights and actions (Aalst 2016). Process technologies combine decision-making probabilities with temporal measures. Machine readable

process models of complex natural phenomena enable numerous applications. As the name implies, process mining AIs "mine" data and surface (e.g. Markov or Bayesian) models of decision-making processes from various input formats. Process mining is also a human-understandable, human-verifiable, and human-explainable AI.

Any structured or unstructured data sources which chronicle events is usable--cyber security logs, Internet of Things networks, information systems, political speeches, social data, policy documents, etc. Three pieces of information are needed to discover processes; additional features enrichen the analysis.

- Case ID: An identifier that represents a specific execution of a process.
- Activity: One of several steps performed within a process. For example, web-based eCommerce process activities might include "Add Item to Cart" and "Initiate Checkout."
- Time Stamp: This orders the activities within each case and enables sophisticated modeling.

CASE	ACTIVITY
1	A
2	A
1	B
1	C
3	A
2	C
3	B
2	B
1	D
2	D
2	E
3	C
3	D
1	E
3	D
3	E

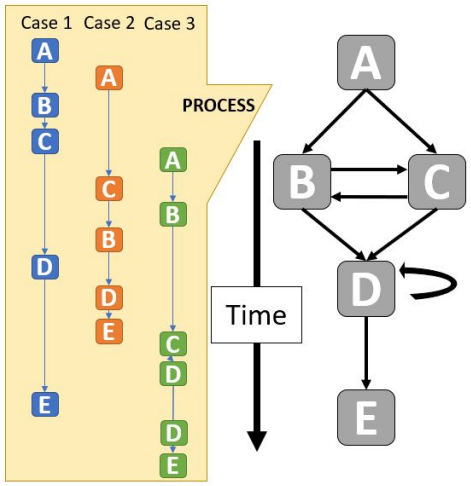

Figure 1

Figure 1 contains a trivial example. Process cases 1 through 3 all contain the same activities, labeled "A" through "E." In Case 1, the activities happen in natural order. In Case 2, Activity "C" precedes "B." Finally, in Case 3, Activity "D" is repeated before concluding with Activity "E." The discovered process accounts for these process variations. Real world processes are significantly more complicated.

Transition Probability Matrix State 1 → 2

State 1 \ State 2	A	B	C	D	E
A		0.67	0.33		
B			0.67	0.33	
C		0.33		0.67	
D				0.25	0.75
E					

Figure 2

Process models are temporal representations of events organized by process case. The model presented in Figure 1 may also be represented algebraically as a Markov transition probability matrix, Figure 2. The flow of information from beginning process activity states, labeled "State 1" contained in the matrix rows, to ending states, "State 2" matrix columns, are expressed as probabilities. Referencing the same simple event log contained in Figure 1, information flows from Activity "A" to Activity "B" two-thirds of the time, and from Activity "A" to Activity "C" one-third of the time. When starting at Activity "D," there is a 25% probability that Activity "D" repeats and a 75% likelihood that Activity "E" follows

Activity “D.” More Cowbell Unlimited’s software uses a heatmap color coding scheme to visualize relative values.

Process mining and social network analysis are conflated frequently. While both techniques have a strong network component, process mining is inherently temporal, as explained earlier. Process models, therefore, may be analyzed using many of the same centrality and distance measures associated with SNA. Additionally, process models enable temporal measures and downstream hand-over-of-work observations not found typically in SNA.

This analysis used a 2015 timeframe data set of tweets (Alfifi et al. 2019). There were approximately 15 million Arabic language tweets; 10 million were Unverified ISIS tweets, and 5 million were verified ISIS tweets. The tweet date range was January 2015-February 2016.

For the verified ISIS accounts, temporal event logs were constructed using the account Twitter handle as the CaseID, hashtag as the process activity, and tweet timestamp as the timestamp. The top 20 hashtags used within the dataset were identified and placed into a unique process activity. Hashtags which were not in the top 20 were placed into an "Other Hashtag" process activity category. Tweets with no hashtags were assigned a "No Hashtag" process activity category. Hashtags were translated from Arabic into English using Google Translate. Finally, a statistically representative sample equal to 0.5% of the dataset was used for process mining.

Results

Two datasets were mined in order to discover latent processes contained within the data. One dataset included the "No Hashtag" and "Other Hashtag" activities, and one dataset removed these activities from the sample. The run which included the "No Hashtag" and "Other Hashtag" activities contained 394 process cases; the sample which removed all of the tweets with “No Hashtag” and “Other Hashtag” has a smaller number of process cases (224 cases) because there were over 150 accounts in the sample which were completely comprised of “No/Other Hashtags.”

More Cowbell Unlimited’s software auto-generates a narrative summary of the processes:

> **Process which Includes “No Hashtag” & “Other Hashtag”**: This process model was created by analyzing 394 cases. The average case time is 13.34 days, and the median case time is 21.81 hours. The min case time is 0.00 seconds. The max case time is 298.12 days. It appears that the case durations are extreme value skewed to the right. There were 46534 total events processed. The activities which were estimated to take the longest amount of time are "No Hashtag" at 1.01 days and "support" at 17.02 hours. There were an average of 118.11 events per case, and a median of 22.0 events per case. There were a

max of 7657 events per case and a min of 1. The most frequent activities are "Other Hashtag" at 5151 sojourns and "No Hashtag" at 5151 sojourns. The process activities with the highest capacity estimates are "No Hashtag" and "Other Hashtag."

Process with “No Hashtag” & “Other Hashtag” Removed: This process model was created by analyzing 224 cases. The average case time is 7.21 days, and the median case time is 10.81 hours. The min case time is 0.00 seconds. The max case time is 204.16 days. It appears that the case durations are log normal skewed to the right. There were 6980 total events processed. The activities which were estimated to take the longest amount of time are "ask forgiveness" at 1.72 days and "support" at 1.61 days. There were an average of 31.16 events per case, and a median of 7.0 events per case. There were a max of 674 events per case and a min of 1. The most frequent activities are "the Islamic state" at 968 sojourns and "News_Change" at 968 sojourns. The process activities with the highest capacity estimates are "the Islamic state" and "the state of conquest."

Figure 3 presents the transition probability matrix obtained from the dataset which includes the “No Hashtag” and “Other Hashtag” process activities. The color coded heat map reveals relatively high probabilities that ISIS Twitter account managers tweet with “No Hashtag” or “Other Hashtag” following a tweet with one of the various top 20 hashtags. Additionally, there is a strong interplay between the “No Hashtag” and “Other Hashtag” process activities. This finding makes sense, as these activities have the highest event counts in the dataset. After tweeting with “No Hashtag,” there is a 76% likelihood that the terrorist propagandist account tweets with one of the many “Other Hashtags.” The “News_Change” hashtag, which might also be interpreted as “News Flash” or “Breaking News,” also has a relatively high information flow into the process activity state. The “praise” hashtag is ignored due to a very small event count

ISIS Twitter Hashtag Transition Probability Matrix

Includes "No Hashtag" & "Other Hashtag" (Number of Cases: 394)

State 2

State 1

Event Node Counts	Anbar Province	Baiji	News_Change	Ninewa	No Hashtag	Other Hashtag	State of Aleppo	State of Sinai	Tikrit	ask forgiveness	gray	praise	storm-packs	support	the Islamic state	the agency _ depths	the state of __Salah_delin	the state of conquest	the state of the pond	urgent	urging
Anbar Province	231			0.39		0.03	0.26	0	0		0.1		0		0.07		0	0.07		0.04	
Baiji	164	0.01		0.5		0.02	0.26			0.02					0.01	0.06	0.14				
News_Change	1177	0	0		0.05	0.01	0.68	0.06	0.01	0	0.02		0		0.01	0	0.09	0	0.03	0.04	
Ninewa	211	0.01	0	0.11		0.07	0.51	0		0					0.12			0.11	0	0.06	
No Hashtag	24313	0.02	0.01	0.07	0.02		0.76	0	0.01	0.01	0.01		0.01	0	0.06	0	0.01	0.01	0	0.01	0
Other Hashtag	15241	0.01	0.01	0.07	0.01	0.39		0.01	0.02	0.01	0	0.02	0.02	0.01	0.2	0.03	0.02	0.08	0.01	0.08	0.02
State of Aleppo	173	0.01	0.01	0.23	0.01	0.16	0.18								0.1		0.01	0.14		0.15	0.01
State of Sinai	135	0.01	0.01	0.09	0.01	0.16	0.16			0.02		0.01	0.01		0.2		0.01	0.21		0.12	
Tikrit	137			0.01		0.28	0.17					0.05	0.04		0.25	0.02	0.09	0.01		0.05	0.02
ask forgiveness	6	0.17				0.33	0.17										0.17				
gray	223	0.05		0.02	0.01	0.19	0.17			0					0.22	0.09	0	0.1		0.14	
praise	1						1														
storm-packs	147	0.02		0.02		0.18	0.35		0.01	0.01				0.01	0.26	0.02	0.01	0.05		0.01	0.04
support	74			0.06		0.49	0.25								0.11					0.08	
the Islamic state	1621	0.01	0	0.03	0.01	0.29	0.34	0	0.01	0.01		0		0		0.01	0.03	0.05	0.01	0.04	0.16
the agency _ depths	190	0.01	0.01	0.18	0.02	0.31	0.34	0.02		0.01		0.01			0.01		0.02	0.02		0.05	0.01
the state of __Salah_delin	361		0.03	0.07	0.01	0.13	0.16	0	0	0		0.01		0	0.03	0		0.16	0.01	0.39	
the state of conquest	706	0.03	0.01	0.02	0.02	0.22	0.43	0.01	0			0.01		0	0.04		0.02		0.03	0.13	0.03
the state of the pond	115	0.01	0.02	0.14	0.01	0.25	0.21			0.01					0.04		0.01	0.01		0.28	0.02
urgent	922	0.03	0.05	0.23	0.01	0.21	0.38	0.01	0.01	0.01		0.01		0	0.02	0	0.02	0	0		0.01
urging	386	0.01	0.01	0.03	0	0.21	0.62	0.01		0.01		0		0	0.1		0.01	0			

Figure 3

.

Figure 4 presents the transition probability matrix obtained from the dataset which removed the “No Hashtag” and “Other Hashtag” process activities. In this result, the “the Islamic state” hashtag is quite pronounced as an inbound activity transition state when compared to the same hashtag state in Figure 1. For example, the “urging” hashtag is followed 67% of the time by the hashtag “the Islamic state.” In Figure 1 the “urging” hashtag transitioned into the “No Hashtag” and “Other Hashtag” states 83% of the time. The “New_Flash” process activity is also more pronounced in Figure 4. Again, the “praise” hashtag is ignored as it has a very small event count.

ISIS Twitter Hashtag Transition Probability Matrix
"No Hashtag" & "Other Hashtag" **Removed** (Number of Process Cases: 224)

State 2

State 1

	Event Node Counts	Anbar Province	Baiji	News_Change	Ninewa	State of Aleppo	State of Sinai	Tikrit	ask forgiveness	gray	praise	storm-packs	support	the Islamic state	the agency _ depths	the state of __Salah_delin	the state of conquest	the state of the pond	urgent	urging
Anbar Province	231		0	0.42	0	0.01	0.02			0.15		0.02		0.13	0.02	0.01	0.1		0.09	0.01
Baiji	164	0.02		0.51				0.03		0.01		0.01		0.05	0.1	0.21	0.03		0.03	
News_Change	1177	0.03	0.03		0.07	0.11	0.06	0		0.05		0.01	0	0.08	0.03	0.15	0.02	0.06	0.28	0.01
Ninewa	211	0.04	0.01	0.22		0.01	0.01	0.02		0.01		0.01	0.01	0.33	0.01	0.01	0.22	0.01	0.09	0.01
State of Aleppo	173	0.02	0.01	0.33	0.01		0.02	0.01		0.02				0.16	0.01	0.03	0.18		0.19	0.01
State of Sinai	135	0.02	0.01	0.18	0.02			0.03		0.02		0.03		0.29	0.01	0.01	0.24	0.01	0.15	
Tikrit	137			0.03		0.01	0.01			0.07		0.05	0.01	0.47	0.03	0.13	0.04	0.01	0.08	0.05
ask forgiveness	6	0.17				0.17		0.17			0.17			0.17		0.17				
gray	223	0.12	0.02	0.06	0.01	0.01		0.01						0.29	0.12	0.01	0.16	0.01	0.17	0.01
praise	1													1						
storm-packs	147	0.04	0.02	0.11	0.01		0.01	0.04		0.02			0.03	0.44	0.03	0.03	0.11	0.01	0.03	0.05
support	74			0.22	0.03		0.03	0.03				0.05		0.32	0.03		0.03		0.24	0.03
the Islamic state	1621	0.05	0.02	0.11	0.04	0.02	0.03	0.03	0	0.03		0.03	0.01		0.02	0.08	0.14	0.03	0.1	0.27
the agency _ depths	190	0.04	0.04	0.3	0.04	0.02	0.03	0.02		0.05		0.02		0.19		0.04	0.11	0.01	0.09	0.01
the state of __Salah_delin	361	0.01	0.05	0.1	0.01	0.01	0.01	0.03		0.01		0.01		0.11	0.01		0.2	0.01	0.43	0.01
the state of conquest	706	0.06	0.02	0.05	0.05	0.04	0.03	0.03	0	0.04		0.02	0.01	0.26	0.04	0.06		0.04	0.2	0.06
the state of the pond	115	0.03	0.02	0.25	0.04	0.01	0.01	0.02		0.01			0.02	0.17	0.01	0.05	0.04		0.32	0.02
urgent	922	0.07	0.08	0.37	0.02	0.01	0.02	0.03	0	0.03		0.01	0.02	0.16	0.03	0.04	0.08	0.01		0.02
urging	386	0.02	0.01	0.07	0.02	0.01	0.01	0.02		0.01		0.04	0.01	0.67	0.02	0.01	0.06	0	0.02	

Figure 4

Figure 5 presents the same tweet data which removed the “No Hashtag” and “Other Hashtag” process activities in a social network representation. This alternative view of the temporal flow of information through the ecosystem keeps only the top 20% of transitions by event count, As with Figure 4, “The Islamic State” process activity is more pronounced due to the relatively high number of entering and exiting tweet events. To a lesser extent, the “News_Change” and “Urgent” hashtag states are also favored by ISIS Twitter account managers.

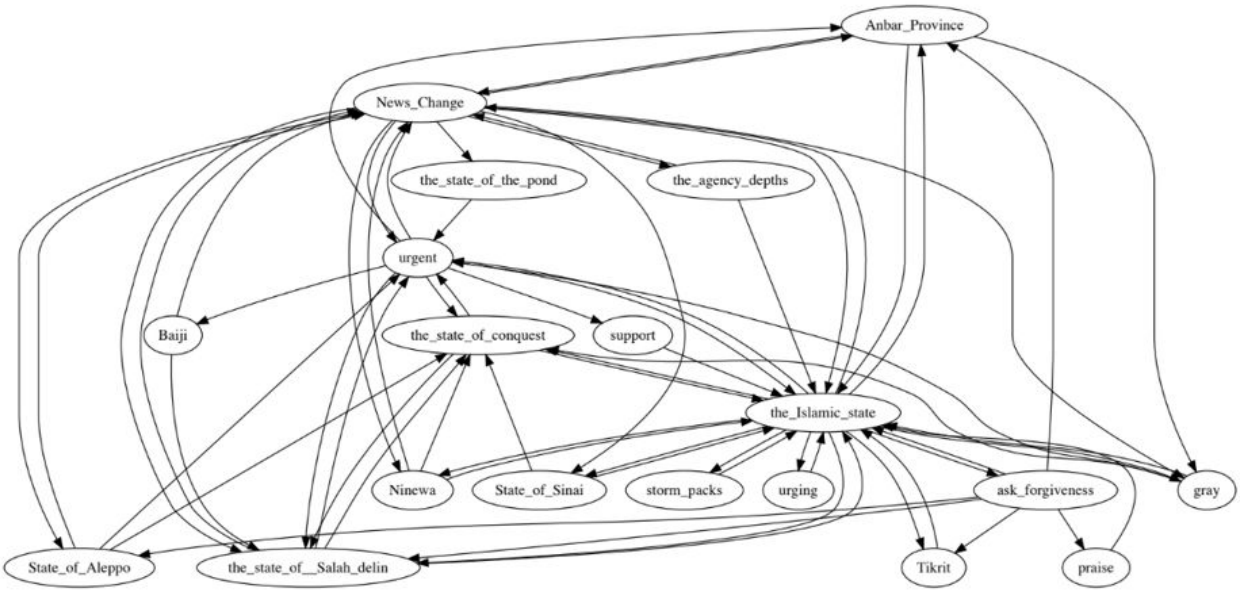

Figure 5

Discussion

Results reveal weaponized bot activity using human-understandable, human-verifiable, and human-explainable AI. ISIS twitter account managers are likely managing dozens or even hundreds of accounts simultaneously with bots. Referring again to Figure 3, after tweeting with an "urging" hashtag, there is an 83% probability the ISIS account tweets with "No Hashtag" or "Other Hashtag." The Markov transition probability matrix simplifies after eliminating those non-propaganda tweets ("No Hashtag" or "Other Hashtag"), and patterns become more pronounced.

A derived hidden Markov model combines the results in Figures 3 and 4 and illuminates the bot logic, which would otherwise be difficult to detect. For example, there appears to be a "hidden state" in the logic of the bot, where it tweets out an "urging" hashtag and then goes into a hidden state where it tweets out random hashtags. It then leaves that hidden state and, with 67% probability, tweets out "the Islamic State" hashtag.

These results are supported by literature and media which discuss propaganda techniques. World War II Nazi propagandists created dozens of benign or innocuous entertainment films in order to camoflage the propaganda intent of the state-sponsored film and media industry while luring the wartime German citizenry to theaters. The propaganda content was delivered to packed theaters as a separate small film or newsreel at the beginning of the feature length films

(*Scorched Earth: Propaganda* 2006). Similarly, Soviet-era RC tactics techniques and procedures (TTP) mixed at least 90% camouflage or background material to blend into the information environment and disguise the true nature of the propaganda (Chotikul 1986). ISIS Twitter account managers, at least in the 2015 timeframe, appear to have adapted these Nazi and Soviet TTPs into their modern social media weaponized bot propaganda campaigns.

Future Effort

This project is a promising start at understanding complex multi-mix choreographed botnet activity and human/bot logic. Detecting malicious social accounts and modeling botnet dynamics in close to real-time is an active area of interest.

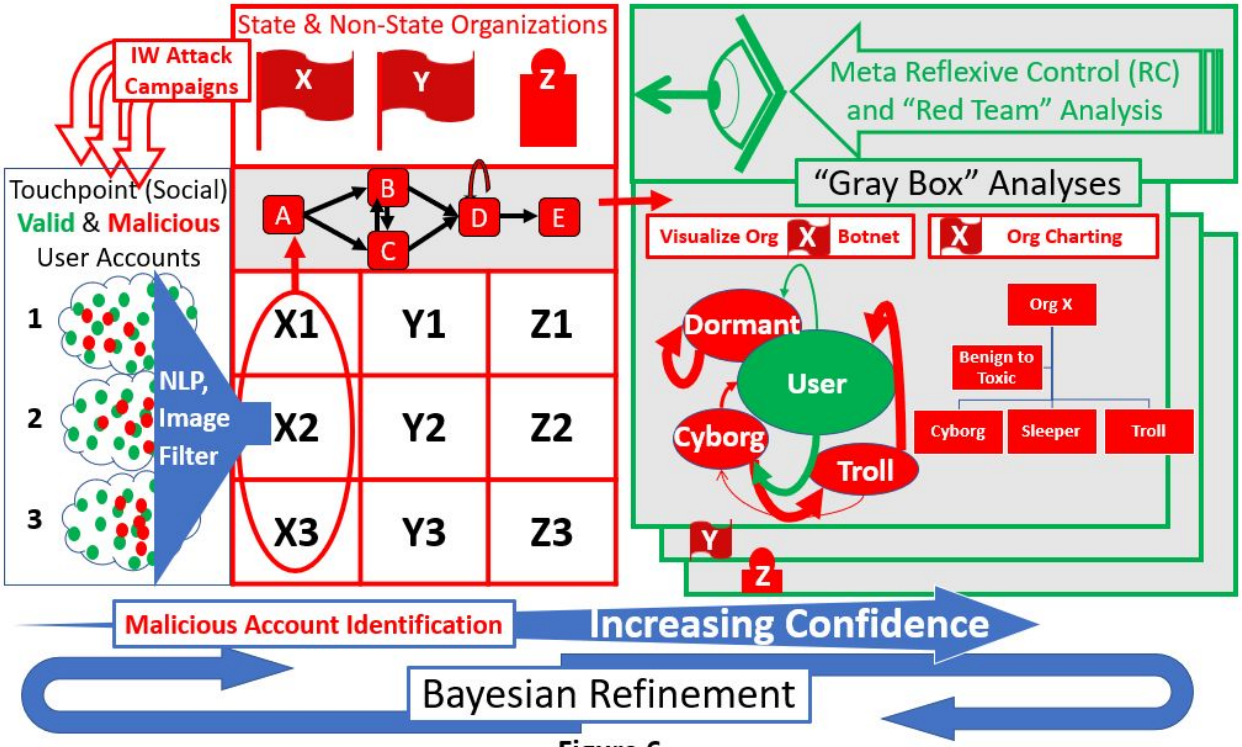

Figure 6

Future efforts should string together creatively several best of breed techniques into a composite classification algorithm in order to continually improve the discovery of malicious accounts and understand choreographed weaponized botnet dynamics which span social media platforms, Figure 6.

Malicious state and non-state actors, labeled "X," "Y," and "Z," inject nuanced, coordinated IW attack vectors into numerous touchpoint platforms such as Twitter, YouTube, WeChat, 8chan.

For simplicity, three touchpoints are displayed and are labeled “1,” “2,” and “3.” IW attack vectors are designed to influence users in ways that benefit the attacker. Based upon recent conversations with experts, Russia’s 2016 election attacks were a “spaghetti against the wall” style of attack just to see what sticks. However, “Russia learned from this experience and is crafting more sophisticated attacks (Cyber Endeavour Conference with non-attributional Chatham House Rules 2019).” Within each touchpoint ecosystem, there are regular user accounts shown as green dots and malicious accounts, which are red dots. The malicious accounts are trolls, bots, cyborgs, sleepers and others which coordinate and proliferate choreographed botnet IW attacks. Assuming an information arms race, adversaries are conducting ever more sophisticated analyses of our society in order to craft powerful, customized, and insidious RC IW attack vectors; additionally, they are developing equally advanced delivery methodologies in order to maximize effect.

Continuing with Figure 6, algorithms and TTPs for detecting and elucidating varieties of organizational accounts of interest and coordinated botnet influence activities from noisy touchpoint data. After ingesting massive amounts of touchpoint data, heuristics are developed to winnow (“filter”) these data to a manageable size. Heterogeneous touchpoint data are combined using AI adapters. Conglomerate process models enable creative temporal measures and reveal coordinated media mix activities, while a Bayesian framework increases malicious account identification confidence. Since ISIS is a non-state group (Organization “Z”), the tweet analysis presented in this paper resides within a single cell of matrix column “Z,” for example cell “Z1.”

Importantly, observe in Figure 6 how malicious IW attacks are likely coordinated across touchpoints and may be represented by an NxN activity matrix. In this model, a 3x3 matrix represents the IW battlespace where each column is a promulgator of IW attack and each row is a touchpoint social platform. For example, State “X” delivers complementary and reinforcing attack messaging vertically across touchpoints 1, 2, and 3. Each cell contains temporal event data which may be mined for processes latent within the data. Enemy organizational inferences and models are possible from contextual process models, as well.

Ultimately, recursive-RC analyses are possible which use process mining to model the enemy’s information warfare efforts. A significant portion of troll, bot, and cyborg processes take the form of social media or are otherwise retrievable off blogs and websites. So, by the nature of their work, nearly all of these processes are overt and available for public view, and thus available for process mining and discovery. Because known adversaries engaged in IW already think of their operations in terms of process modeling, process mining AI is a framework especially well-suited to analyzing their activities and modeling their behavior, whether as direct RC or within a meta-RC or recursive RC-context. Such an analysis may also be thought of as a “red team” examination which illuminates fault lines in societal discourse. Red Team projects are

intensive self-examinations from an adversary's perspective and are a well-known technique to build up defensive capabilities by having friendly forces simulate enemy tactics to discover weaknesses (Zenko 2015).

Models derived from process mining using domain-specific NLP detect botnet behavior around 'hot button' topics of strategic or tactical importance, with process mining of enemy behavior helping to elucidate what those areas are. Known adversaries already think of their IW operations in terms of process modeling; therefore, process mining AI is a framework especially well-suited to analyzing their activities and modeling their behavior. Red team examinations catalogue and monitor continually fault lines in social discourse for IW intrusion. The potential exists to develop sophisticated IW defense tools, real-time attack detection techniques, and business certifications that protect organizations and larger society from IW well beyond social media (Bicknell and Krebs 2019a; 2019b).

Conclusion

Working with Texas A&M's Cyber Security Center, weaponized AI information warfare attack vector botnet dynamics were discovered. Using a 2015-era ISIS Twitter data set, bots appear stateful between tweets and set up a sequence of tweets over time with a structured logic. In other words, ISIS Twitter bots are not just randomly tweeting; however, some of the deliberate logic may indeed be random filler tweets which have no hashtags ("No Hashtag") or lower count hashtags ("Other Hashtag").

These results are pioneering and suggest two interesting findings. First, evidence of bot activity which would otherwise be difficult to detect was discovered using process technologies. Moreover, the results suggest ISIS twitter account managers deploy information camouflage as part of their choreographed information warfare campaigns similar to World War II Nazi propagandists and Soviet-era reflexive control practitioners. These results warrant further investigation and should be incorporated into future social media botnet discovery and modeling efforts.

Adapting this technique into the corporate or government cyber environment may provide new perspectives relevant to the cyber fight, as well. For example, blue team cyber security defenders may more readily connect and understand disparate attack vectors as part of choreographed campaigns (Bicknell and Krebs 2019a). Similarly, satellite networks are likely amenable to information camouflage and hidden attack vector discovery. Finally, geographic areas of interest may be analyzed in a process framework to discover patterns of life deviations--including deliberately camouflaged illicit activities ("Large Data Aggregation from Small Satellites to Determine Patterns of Life Modifications" 2019).

Future research and development should advance hidden Markov model (HMM) and Long short-term memory (LTSM) techniques for discovering important temporal relationships. User feedback will inform improvements to process visualizations. When persistent use cases are discovered, “gray box" analyses, which require expert assistance and model interpretation, may be converted into close to real time black box AI for discovering IW attacks in close to real time.

Citations

Aalst, Wil M. P. van der. 2016. *Process Mining: Data Science in Action*. 2nd ed. 2016 edition. New York, NY: Springer.

Alfifi, Majid, Parisa Kaghazgaran, James Caverlee, and Fred Morstatter. 2019. “A Large-Scale Study of ISIS Social Media Strategy: Community Size, Collective Influence, and Behavioral Impact.” *Proceedings of the International AAAI Conference on Web and Social Media* 13 (July): 58–67.

Bicknell, John W, and Werner G Krebs. 2019a. “Process Dominance: The Capability Nobody Is Talking About.” 2019. https://morecowbellunlimited.com/wp-content/uploads/Process-Mining-in-DoD-Context-White-Paper.pdf.

———. 2019b. “FOCAL Information Warfare Defense Standard.” ResearchGate. June 14, 2019. https://www.researchgate.net/publication/333774135_FOCAL_Information_Warfare_Defense_Standard_TM?channel=doi&linkId=5d03306b4585157d15a95823&showFulltext=true.

Chotikul, Diane. 1986. “The Soviet Theory of Reflexive Control In...” NOS55-86–013. Monterey, California: Naval Postgraduate School. http://nsarchive.gwu.edu/dc.html?doc=3901091-Diane-Chotikul-The-Soviet-Theory-of-Reflexive.

Cyber Endeavour Conference with non-attributional Chatham House Rules. 2019.

“DoD Early Detection of Information Campaigns by Adversarial State and Non-State Actors.” 2019. 2019. https://sbir.defensebusiness.org/topics?topicId=30637.

Iyengar, Rishi. 2018. “WhatsApp Has Been Linked to Lynchings in India. Facebook Is Trying to Contain the Crisis.” CNN. July 27, 2018. https://www.cnn.com/2018/09/30/tech/facebook-whatsapp-india-misinformation/index.html.

“Large Data Aggregation from Small Satellites to Determine Patterns of Life Modifications.” 2019. 2019. https://hyperspacechallenge.com/large-data-aggregation-from-small-satellites-to-determine-pattern-of-life-modifications/.

McLaughlin, Timothy. 2018. “How Facebook’s Rise Fueled Chaos and Confusion in Myanmar.” *Wired*, July 6, 2018. https://www.wired.com/story/how-facebooks-rise-fueled-chaos-and-confusion-in-myanmar/.

Novikov, Dmitry A., and Alexander G. Chkhartishvili. 2014. *Reflexion and Control : Mathematical*

Models. CRC Press. https://doi.org/10.1201/b16625.
Paul, Christopher, and Miriam Matthews. 2016. “The Russian ‘Firehose of Falsehood’ Propaganda Model.” Product Page. 2016. https://www.rand.org/pubs/perspectives/PE198.html.
Scorched Earth: Propaganda. 2006. https://www.amazon.com/Scorched-Earth-Propaganda-Unavailable/dp/B01DEFZUMA.
Singer, PW, and Emerson T Brooking. 2018. *LikeWar: The Weaponization of Social Media*. Eamon Dolan/Houghton Mifflin Harcourt. https://www.amazon.com/LikeWar-Weaponization-P-W-Singer/dp/1328695743.
Thomas, Timothy. 2004. “Russia’s Reflexive Control Theory and the Military.” *The Journal of Slavic Military Studies* 17 (2): 237–56. https://doi.org/10.1080/13518040490450529.
Waltzman, Rand. 2017. “SASC Testimony: The Weaponization of Information.” Product Page. https://www.rand.org/pubs/testimonies/CT473.html.
Whitehead, Alfred North. 1979. *Process and Reality*. 2nd edition. New York: Free Press.
Zenko, Micah. 2015. *Red Team: How to Succeed By Thinking Like the Enemy*. 1 edition. New York: Basic Books.