

Propaganda Filters: Tracking Malign Foreign Interventions on Social Media

Bruce Forrester

Defence Research and Development Canada, Quebec City
Canada

bruce.forrester@drdc-rddc.gc.ca

Akiva Bacovcin

Nexalogy, Montréal, Canada

abacovcin@nexalogy.com

Zachary Devereaux

Nexalogy, Montréal, Canada

zdevereaux@nexalogy.com

Stefany Bedoya

Nexalogy, Montréal, Canada

sbedoya@nexalogy.com

ABSTRACT

Democracy requires that the people have access to factual information and that they be able to participate in open and honest discussion on issues of importance. Unfortunately, there are forces in this world that would benefit from weak democracies. Many such forces are currently very active on social media spreading propaganda, disinformation, and trying to increase divisions within democratic populations. Identifying foreign-backed users and BOTs that are being exploited to interfere in democratic processes is an important first step to understanding effects. This research aimed to describe a methodology that can be used to identify suspicious users from which indicators of foreign narratives can be inferred. To help describe the network effects of BOTS, the BEND (build, engage, neutralize, distract) framework is used. To make the research both timely and real, data collection and analysis was conducted on divisive issues found within Canada. We identified evidence of BOTs, with attribution to Russia, meddling in Twitter discussions on these wedge issues. This meddling was mapped onto the BEND framework. In particular, this meddling seemed to be focused in Boosting and Enhancing extreme viewpoints.

1.0 INTRODUCTION

Evidence of foreign interference in numerous democratic processes, perpetrated using social media, is no longer surprising [1-4]. In fact, Bradshaw and Howard [5] have found that in at least 70 countries there is evidence of social media manipulation to shape public attitudes. NATO is also well aware of the 2013 “Gerasimov Doctrine”, Russian Chief of the General Staff that promotes indirect and asymmetric methods [6] that uses the information environment to attack human cognition centering on distraction and manipulation [7]. As an example of manipulation, according to a NATO study [1], Russia had four strategic (read political) goals for their propaganda effort against the Ukraine:

1. To promote Russia as a crucial player in the polycentric world in the process of international peace;
2. To claim Russia’s superiority over the US;
3. To prevent Ukraine’s transformation into being part of the external border of NATO and the European union; and,
4. To soften and, in the nearest future, achieve the lifting of the sanctions regime against Russia.

Within Canada, the focus is more about distracting the population. Upsetting well-established democracies by increasing the divisions between citizens’ with opposing views is an effective method; while the people of that country are busy “fighting” each other, Russia is able to move with greater freedom with less scrutiny. This research was conducted with regard to the several of the objectives of the IST-177 research task group. In particular:

1. Exploration of attack vectors and effects of social media influence attacks;
2. Fake News and Propaganda Recognition;
3. Understand tools for injection of military-relevant information into social media streams;
4. Exploration of relevant technologies and tools; and,
5. Analytics for segmenting target populations for maximum influence opportunity.

The goal of this research was to try to detect Russian influence within divisive Canadian issues. Additionally, we hoped to decrease analyst time to insight for these issues by using AI to quickly treat large amounts of data e.g. big data and automation.

In order to find evidence of Russian interference on social media, analysts’ must understand the likely target areas within the information environment. They must collect, sort, and analyse hundreds of thousands of social media posts. A daunting challenge within the world of data analysis, for situational awareness, public affairs, and numerous other purposes, is the amount of time it takes to analyse enormous groups of social media posts. Indeed, in many cases, there is not enough manpower available to make manual inspection possible in a useful timeframe. This research aimed to help address this challenge by using filters and AI based treatment of big data via neural networks which have been developed to significantly decrease the time to insight with large data sets of social media.

The novel methodology derived here is based on knowledge of Russian (Maskirovka) techniques. Their information warfare tactics focus on promoting division in foreign populations by finding wedge issues to push people to the extremities. In studying the behaviour of Russian BOTs on social media, McKelvey and Dubois found “amplifiers, dampeners, transparency, and servant BOTs” [8, 9].

In order to better understand the motivation and tactics of the identified BOTs, the BEND framework [10] was used. This framework is based on research on disinformation and Russian information manoeuvres and describes strategies that can be used by actors in the information environment. Information activities are characterised along two axis: community manipulation versus content manipulation and positive manipulation versus negative manipulation. Each letter of BEND corresponds to one quadrant produced by

the combination of the two axis Figure 1. For example, the positive community manoeuvre ‘Build’ helps to create a group of users that supports a brand, topic, or theme. A negative content manoeuvre such as ‘Distort’ manipulates content through omitting facts, disinformation, or lies, in order to alter an audience’s perception. For a full discussion of this framework and description of the 16 manoeuvres refer to [10].

The BEND Framework

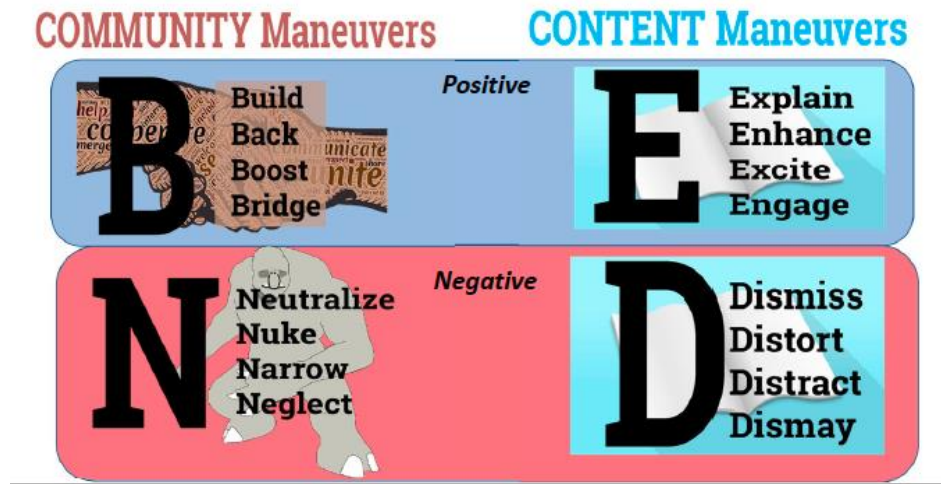


Figure 1. The BEND framework

2.0 METHODOLOGY

Keeping in mind the research objectives identified in the introduction, a repeatable and simple method was devised that allowed for iteration over the data and spiral analysis of interim results between steps. We were interested in determining if there was evidence of Russian interference in Canadian based Twitter discussions that dealt with divisive issues. We assumed that Russian information warfare tactics would be centred on the further trying to divide (encouraging more extreme views and opinions) the Canadian population by using community and content manoeuvres as identified by the BEND model. As such, the first step was to determine topics that were potentially divisive. These issues were chosen based on their importance to many Canadians and because there exists strong opposing opinions within the issues, often based on ones’ personal beliefs.

A recursive and iterative, discovery method was used for this research, which was applied to several different, but related, datasets. The overall method used for this research was as follows:

1. Identify possible wedge issues for foreign interference and related search terms;
2. Perform initial discovery & filtering;
3. Apply predetermined filters and Authentic Chatter map to detect suspicious users;
4. Analyse initial results;
5. Identify and classify users;
6. Update filters to include suspicious users (likely to be BOTs);
7. Apply updated filters to detect associated users and related narratives;
8. Identify network effects according to BEND; and,
9. Apply updated filters to examine inter-issue narratives.

Starting in June 2019, we identified a number of potential wedge issues based on current media political discussion and our own knowledge of the Canadian political scene. Those issues were: immigration, ethical issues, bill21, Canada-China relations, Canada-USA relations, climate change, pharmacare, indigenous issues, and pipelines. A control group using an overarching hashtag was also collected to permit discovery of unidentified issues.

Initial data collect based on the Twitter API (about 1500 tweets) was performed and the results were examined. Boolean search queries were created to capture the relevant terms, checking the initial version for accuracy, and then performing necessary edits until the data capture was accurate and comprehensive. Top Hashtags and Top Words, for example, contributed to search strategy refinement and data veracity was achieved by iterating through this process. Contradicting hashtags, on either side of an issue, was evidence of a wedge issue.

Data collection occurred from 11 June to 10 Aug 2019. For this short paper, the research investigate only three topics: ethical issues, pipelines, and climate change. It is always easier to capture a wider swath around a topic and then filter and refine the results. Stop words included “Trump” and geolocations derived from Twitter based on latitude and longitude or on proprietary location derivatives (likely biographical, discussion topics) were used to increase the probability that tweets were coming from Canada. The goal was to create Canadian specific data collections without, as much as possible, noise from the USA.

The rest of this section will be a walkthrough of the methodology. The method was very iterative, often examining the summary ‘reports’ provided by the NexaIntelligence software. The software used allowed for quick examination of the overall content through these numerous reports that were automatically produced. Ranks and clustering of the data facilitated the understanding of overall themes and groupings. Several reports (automatically produced by the NexaIntelligence software) are presented in the results section.

Starting with a chosen data set (ethical issues which was just under 30k tweets) tweets were examined to see if they contained issues that would be a target of Russian interference, such as anti-western rhetoric. Next, two filters were applied to the dataset. The first filter consists of users that have shared Russian propaganda content appearing on a list of 200 websites taken from www.propornot.com. This site has, and continues to, identify sites that produce or propagate Russian propaganda. The second filter identifies authors who are associated with the Internet Research Agency (IRA) and who have been labelled as Russian Trolls by Twitter. This data set was released by Twitter in 2018. If the filters produce results, the tops users were scrutinised (see next paragraph) and the tweets’ contents were further examined for extreme narratives. What content might support the Russian agenda? To whom would a group of BOTs, promoting such a view, be beneficial?

The next step took a closer look at the users identified by the filters. Users were examined to determine the nature of their account (i.e. is it a BOT? Where are they from?). A prototype analysis tool was applied Figure 2, based on Authentic Chatter [11], that classifies users as either BOT or Official (an official NATO, Government or organizational account) using machine-learning. As well, the users’ Twitter sites were qualitatively examined to confirm the automated findings. The following steps were used to verify that a suspected BOT, Figure 3, was actually a BOT after detection by the Authentic Chatter method:

1. Check if user exists;
2. Look at user photo;
3. Look at user name for oddities;
4. Look at number of tweets compared to date of creation;
5. Look at following vs. followers (are they similar?);
6. Check for posting behaviour e.g. all retweets?
7. Look at likes;
8. Use common sense regarding naming e.g. ‘nationalist’;

9. Use common sense regarding the bio; and,
10. Reverse image lookup on suspicious profile pictures and do a Google search.

Once suspicious users were classified as BOT, a new set of filters were developed to include these BOTS. As additional BOTS, or users that interacted with the BOTS, were found in the various datasets, these too were added to the filter. The new filters were then run using previously examined datasets. Again content derived from this new filtering was examined and compared to the previous results.

The above steps were designed to filter the dataset down to the suspected Russian content. At this point the content was again examined as were the top hashtags and top words. The user interactions were examined to look for network effects; was there evidence of amplification or any of the BEND manoeuvres?

The dataset was then filtered to include only the users which were highly suspected of being Russian controlled or users that were interacting with the BOTS. The content was examined for narrative imbalances. Finally, the three data sets were cross-examined for inter-narrative similarities and for BOTS acting across wedge issues.

3.0 RESULTS

3.1 Ethical issues Dataset

The period of analysis was from the 11th of June to the 10th of August. In total, there were 7,174 results from 3,694 publishers (users). The initial attempt to identify Russian propaganda in this dataset used pre-set filters (blacklists) to look for examples of tweets produced by known sources of Russian propaganda (again, derived from the Twitter IRA dataset and based on a previous blacklist constructed from users that frequently link to Russian propaganda sites identified by PropOrNot). In this dataset, there were no examples of data from or mentioning users from either of these blacklists.

Since the pre-existing blacklists failed to identify and users, a prototype tool developed for identifying different categories of users on Twitter was employed: the Authentic Chatter chart, Figure 2. This tool relies on machine learning algorithms trained on examples of official accounts (media personalities, government agencies, corporate accounts) and BOT accounts (as identified by Botometer, BotSentinel and botcheck.me). It gives a probability of any account being an official account, a BOT, both or neither. Attention was focused on accounts with more than a 50% chance of being a BOT.




Figure 2. Automatic detection of Official or BOT users. The users in the bottom left quadrant are likely to be ‘authentic users’ not controlled by others not part of an official organisation.

The result of this filtering gave 43 actors that have at least 50% probability of being a BOT or Troll, producing 316 posts out of 23,819 overall collected results. Hashtags from these accounts are strongly negative towards a left of centre viewpoint. Some of the most active accounts, including @Working_Cdn2019 have (as of September 2019) been taken down although the Nexalogy software still preserves the data that was collected from Twitter before an account suspension. Another suspected BOT was LisaMar91564392, Figure 3. This account was examined for BOT-like qualities.



Figure 3 The Twitter page for user LisaMar91564392

Looking at the content of the suspected BOTs, it was evident that they focused on a right leaning stance within ethical issues. The BOT’s narrative focus can be seen by (a) the list of hashtags used, and (b) by manually examining the 75 tweets from the set (as this was such a small number of tweets, manual examination was feasible). While a set of automated accounts that were pushing a pro-Russian narrative was

identified, the next step was to attempt to determine if there is a foreign connection to these BOTs.

3.2 Pipelines Dataset

The pipeline dataset contained 80,903 posts and 57,620 publishers from June 11th to August 10th. To look for evidence of Russian interference, the propaganda filters were again applied. While the IRA list returned no results, the PropOrNot list found 83 posts by 73 users, confirming that at least some of the users were associated with Russian propaganda sites in this Canadian pipeline discussion. Investigating the top hashtags, Figure 4, and manually inspecting a sample of tweets indicated that these users were pushing a largely leftward (politically) leaning set of themes e.g. destroying organic farms or pro-Iranian and anti-Saudi messages. There were several discussions of Cold War themes that seemed more suspicious given the original topic of discussion.

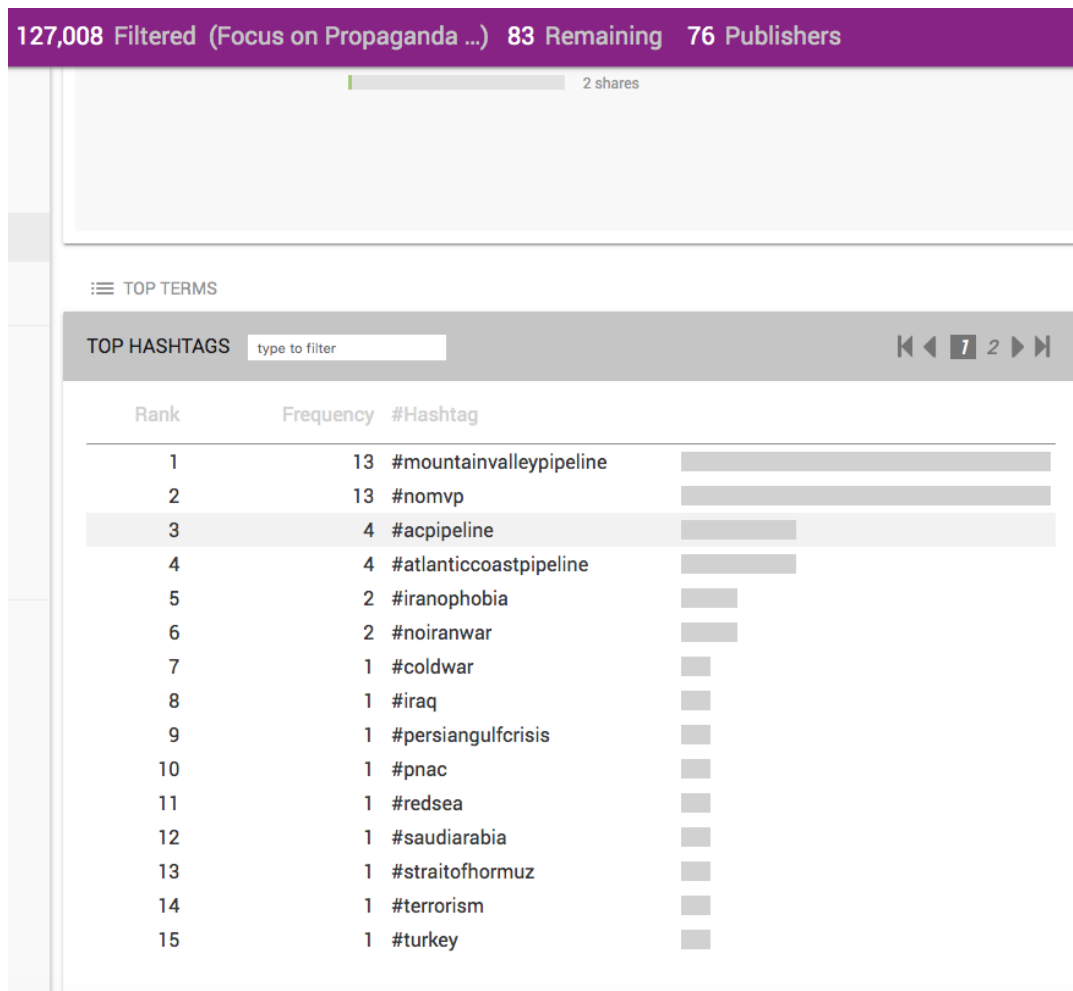


Figure 4. The top hashtags associated with the 73 users identified using the Russian propaganda content filter.

In order to examine if the BOTs from the ethical issues dataset also occurred in the Pipeline dataset, a search for tweets that mentioned or were produced by the same set of BOTs from ethical issues was conducted. Taking the top 15 most prolific BOTs from ethical issues, a filter was made and applied to the pipeline data for these users. The result found 22 posts and 13 users of which 5 users were the same. All of the 8 (13-5) new users were examined. An interesting tweet was found, Figure 5. It contained a long list of '@users'

(with sort messages and link at the end - this link leads to a version of God bless the USA by Lee Greenwood) that looked suspiciously BOT-like as well. Upon verification by searching for each users' Twitter page, it was found that many of these suspicious users have been deleted by Twitter.

```
. description: "@BrentDgls @Danbalkwill1 @BobRae14 @georgewcowles
@MonicaFibonacci @Canadian_Chris_ @lassiter1550 @OldManDuke @muchmore2cents
@pacopilbakalao @phil_rack @leighgt @indivcan @backmarker3 @JohnAll63815927
@grossw04 @MarkMe60 @GirouxRheal @trains72 @torontoensis @Blueyes9445
@vesnalaaurie8 @HaveWeAllGoneM1 @NewImproved9 @jay_slatter @daveb2561
@AsuAdanac @defiantcanuck @Friends0Science @Canadian_logic_ @IsExtortion
@RRidley11 @LDawg05 @kstokesvies @SusanIverach @lambert_pp @JIsaacsonV2
@laShawner951 @nemo_gratis @FayMary3 @jacksurfs204 @UpKeeks @molly6342
@Mim19561 @ThedeplorableM @DawsonMdhoust @JonDavi87753992 Sounds just like
the Pipeline expansion ??? https://youtu.be/Q65KZIqay4E"
```

Figure 5. A suspicious tweet that has a large sender list and a short message with a link.

Having found this interesting connection and its corresponding tweet, the original 15 users from the ethical issues set and the @user mentions from the suspicious BOTs were combined to create a new Ru specific blacklist filter. In order to see if there was a connection between this blacklist and Russian propaganda, the intersection between the Ru blacklist and the PropOrNot blacklist within the pipeline data was examined. The results, showed 50 users making 75 tweets for the time period. This showed overlap in terms of BOT and suspicious activity between the Canadian pipeline and the ethical issues discussions.

3.2.1 BOT Narratives in Pipeline Dataset

To evaluate the BOT narratives, the pipeline dataset was filtered for users or mentions of users present on one of the two blacklists (the new Ru blacklist and the PropOrNot blacklist). There was a high-level discussion of why climate change, as provoked by the Anthropocene, is an incorrect hypothesis. People were actively engage with the @friendsoscience, which was being promoted by the BOT network. This promoting consisted of replying to those whom reply to the account, sharing anti-renewable energy memes, and anti-left messaging. Using the BEND framework, it was found that:

1. DISTORT was present in response to the 12 year window for the point of no return on climate change;
2. AMPLIFY of Albertan pundits such as Friends of Science Barry Cooper was a pattern of behaviour by pro-Russian accounts, perhaps in order to foment Alberta separatism or undermine Canadian federalism? Previous research also showed that Ezra Levant was promoted by Russia-friendly accounts. There is also anti-protester and pro-pipeline narrative shaping. For instance, @LisaMar91564392 was present in this cross-section as an anti-left, but pro-pipeline oil & gas, and this BOT was present in the ethical issues topic issue. BOTs also pushed a Libertarian viewpoint.
3. AMPLIFYING climate change skeptics and far-right politics including anti-left sentiment in Canada is the main theme in this cross-section, in particular identified BOTs are AMPLIFYING @justinshafer1 whose profile links CPC, NRA and US Republicans.

3.3.2 Cross-Section Inter-Narrative Analysis

A filter that combined the BOTs from ethical issues and the Pipeline datasets was now applied to the previously examined ethical issues data. This was testing the attribution of Russian activity derived from the method to this point. There were 13 users with 59 tweets. These users were part of BOT nets discovered in both the ethical issues and Pipeline discussions. Looking at the top hashtags and words, Figure 6, the BOTs' expressed extremist views towards the left. The BOTS that were associated with Russian propaganda in this data sample were anti-left.




Figure 6. Top Hashtags and Words associated with the combined ethical issues and Pipeline BOT filters

3.3 Climate Change Dataset

In order to clarify the understanding from pipelines, the broader wedge issue of Climate Change, containing 500,499 tweets, was used for validation. The original BOT list from the ethical issues dataset was applied to Climate Change resulting in 9 publishers making 16 Tweets. This revealed some new BOTS but also some ‘useful idiots’. Useful idiots are defined as ordinary users who naively and unknowingly follow and spread the messages promoted by BOTs. Next the propaganda content filter revealed 445 tweets from 400 publishers. This filter was combined with the filter from ethical issues and Pipelines, which revealed 227 publishers making 491 Tweets. An actor interaction map, Figure 7, was produced. The actor interaction map shows the social network of conversations and their directionality on Twitter, with Influencers in green that are referenced more than they mention others, and loudspeakers in blue who reference others more often than they are referenced themselves. It shows that there were two tiers of targets in Green, e.g. larger volume influencers and smaller volume influencers within the core malign network. The targeting, it appears, was done by two BOT networks, visible in the upper and lower quadrants of the map. There were BOTs that made tweets mentioning many people and then the targets’ replies to those tweets, and those replies often included the original @ list. This technique effectively amplified the propagation of the message, as the initial list was made by a BOT but the replies were made by real human users. This is a known technique used to game the Twitter algorithm. This is an ENGAGE strategy according to BEND. There was also BOOST in terms of amplifying the voice of climate change skeptics. This is another way to use Nexalogy to discover narratives, for example the top two most mentioned accounts, mentioned by BOTS are climate change sceptics, but ones that have ‘nuanced’ positions against climate change as opposed to simple refusal. What this shows is the DISTORT and DISMISS techniques also being applied.

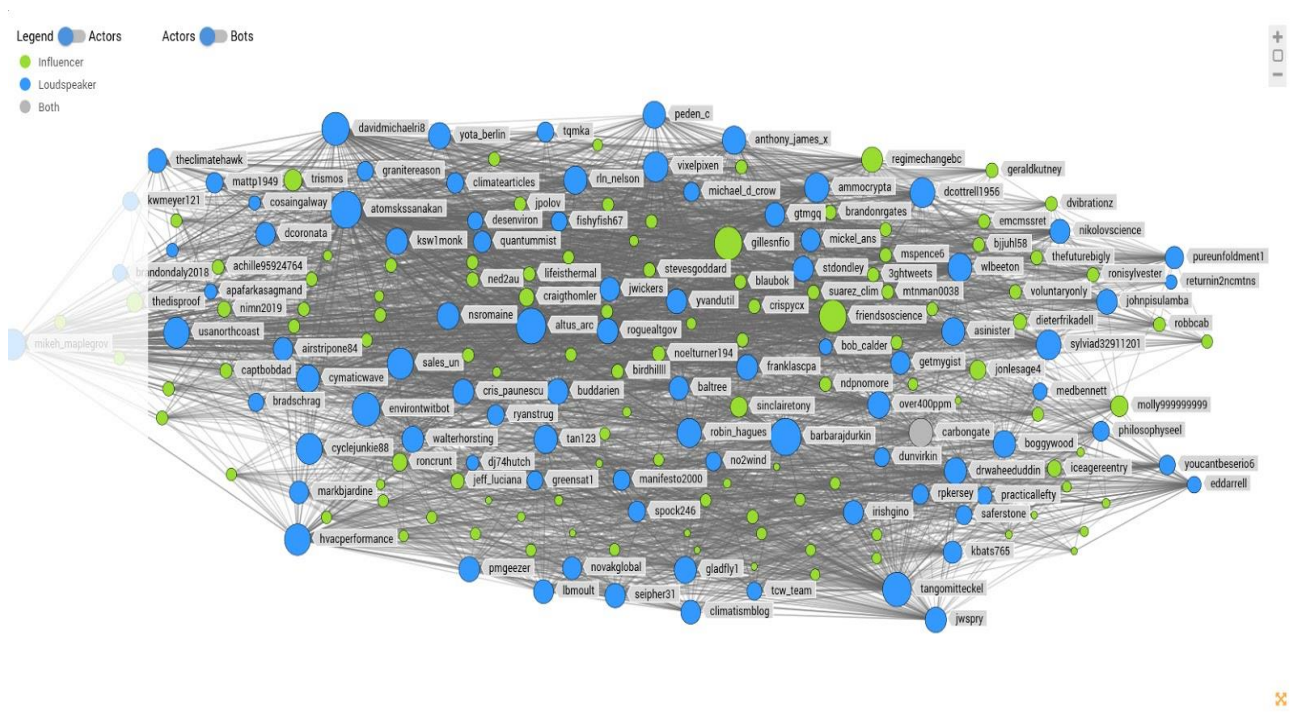


Figure 7. The actor interaction map showing the connections between BOTs identified in the ethical issues and Pipeline datasets applied to the Climate Change dataset.

4.0 DISCUSSION

A set of suspicious users was found using two Russian propaganda filters. A new AI BOT detection technique identified additional suspicious users, some that were identified as overlapping. The narratives pushed by this BOT network emphasised an anti-western, far-right approach. In particular, it demonstrated an anti-western perspective in the ethical issues and in the pipeline topics. In the climate change discussions they the BOTs supported pipelines and opposed the attribution of climate change to human activity. This resembles previous research on the action of Russian agents in the 2016 United States presidential election [12].

As mentioned in the introduction, the BEND information manoeuvres framework, Figure 1, allows analysts to map information activities to tactical categories. In our investigation, we observed the following types of BEND tactics being implemented.

In the Content axis of BEND, ENHANCE was being used in relation to anti-tarsands (the oil producing region of Canada) and anti-western content, and the narratives being promoted involve negative emotions - a DISMAY tactic. In the Community axis of BEND, BOOST and BUILD were observed across wedge issues using the BOT network. The BACK manoeuvre takes place in @mentions of the party leaders. BRIDGE occurred between wedge topic areas. To confirm BEND findings (especially with regards to the Community aspects) a SNA tool like ORA-pro could be used to increase our confidence levels - for example in order to see propagation within the network and compare BOT scoring across tools.

The results from the propaganda filters were used to attribute Russian influence in the results. However, there is a potential of bias stemming from the use of geolocation in the original queries. The benefit of geolocating is less overwhelming data, but the con is that it can bias the dataset towards Twitter’s definition of a location source. Twitter queries involving geolocation take advantage of a combination of Machine

Learning attribution of the author of a tweet to the selected location, and actual geocoding of tweet. Since the searches looked specifically for data from inside Canada, tweets originating from Venezuela or Russia would have been eliminated. Hence, leading to a bias within the sample against foreign influence not pretending to originate within Canada. To verify the degree of bias, the ethical issues dataset was also captured globally as of Sept. 18th and, as it is a very Canadian specific issue, will be checked for potential users from outside of Canada. This post-analysis tweaking reinforces a general methodological point that the queries for any dataset should be re-examined every few months and modified if necessary (to check if bias has been detected or new indicative terms have entered the discussion).

This method can only produce an inferred attribution due to the use of the Russian propaganda filter. This attribution may not be valid. It is very hard to trace exactly who controls a BOT. There were 5 BOTs that overlapped between Pipelines and ethical issues. These two issues are quite different, and one might not expect the users to be similar between topics. This adds to the indication of interference by a third party. It makes sense that Russian BOTs would promote anti-pipeline (#tarsands is the only hashtag in the overlap sub-set) and also promote anti-western discussions, in order to promote Russian geopolitical interests. This makes sense from a Russian perspective because less oil produced by Canada, and being against users seen as pro-pipeline, means less Canadian oil would make it to market.

Many of the BOT accounts that were found in the datasets has since been deleted by Twitter. So, the fact that accounts like @Working_Cdn2019 have been disabled by Twitter shows that there is uptake on Twitter's behalf and actions undertaken in relation to foreign interference within Canadian discussions (or perhaps these same BOTs are being used in other topic areas). Because tweets persist on the NexaIntelligence software, this research was able to detect these now deleted accounts, whereas the removed accounts would not be obvious to everyday users.

5.0 CONCLUSION

We identified evidence of BOTs, with attribution to Russia, meddling in Twitter discussions on Canadian wedge issues. This meddling was mapped onto the BEND framework. In particular, this meddling seemed to be focused in BOOSTING and ENHANCING extreme viewpoints on wedge issues. Focusing on the BOT net identified during this analysis, the propaganda promoted a populist/far-right perspective on establishment politics and pipeline/climate change development (human caused climate change denial).

This research was primarily interested in determining the effectiveness of the method and tools for identifying possible foreign influence. Follow on research will look at suspicious users and narratives and how these change over time. While the methodology will be further streamlined, the current NexaIntelligence tool can already be used in identifying foreign information warfare activity on Twitter. In particular, using a combination of blacklist filters, AI driven tools and analytic methodology, an analyst can quickly detect BOT activity, determine probable attribution, identify key messages and narratives, and determine the tactics using the BEND model. Without these tools, this type of activity would take significantly longer to detect and understand.

One of the current challenges to this methodology is correctly identifying the relevant narratives and then determining the prevalence and development of those narratives across the Twitter community and through time. To enhance narrative analysis, a tool is in development that will detect 'Tweet Caboodles' (individual tweets that are tightly connected in meaning). These caboodles can be combined to identify higher-level issues and topics. Caboodles, issues, and topics will be integrated to help identify how different types of narratives evolve and are propagated by the social network.

6.0 REFERENCES

- [1] NATO, *The dynamics of Russia's information activities against Ukraine during the Syria campaign*. . 2016: NATO Strategic Communications COE, Riga.
- [2] Berzins, J., *Russia's new generation warfare in Ukraine: Implications for Latvian defense policy*, C.f.S.a.S. Research, Editor. 2014, National Defence Academy of Latvia. p. 15.
- [3] Paul, C. and M. Matthews, *The Russian "Firehose of Falsehood" Propaganda Model*. 2016, Rand Corporation.
- [4] Shane, S., *The Fake Americans Russia Created to Influence the Election*, in *The New York Times*. 2017: New York.
- [5] Bradshaw, S. and P.N. Howard, *The Global Disinformation Disorder: 2019 Global Inventory of Organised Social Media Manipulation. Working Paper 2019.2*, in *Project on Computational Propaganda*. 2019: Oxford, UK.
- [6] Bartles, C., *Getting Gerasimov Right*. Military review, 2016. 96: p. 30-38.
- [7] Allen, T.S. and A.J. Moore, *Victory without Casualties: Russia's Information Operations*. Parameters, 2018. 48(1).
- [8] McKelvey, F. and E. Dubois, *Computational Propaganda in Canada: The Use of Political Bots*, in *Computational Propaganda Research Project*. 2017, University of Oxford: Oxford, UK.
- [9] Gorwa, R. and D. Guilbeault, *Unpacking the Social Media Bot: A Typology to Guide Research and Policy*. Policy & Internet, 2018.
- [10] Beskow, D. and K.M. Carley, *Social Cyber Security: An Emerging National Security Requirement*. American Military Review, 2019.
- [11] Forrester, B., *Authentic Chatter*. Computational and Mathematical Organization Theory Journal, In Press.
- [12] Ruck, D.J., et al., *Internet Research Agency Twitter activity predicted 2016 U.S. election polls*. First Monday, 2019. 24, Number 7 - 1 July 2019.

