ARL-TR-8732 ● JULY 2019

DEVCOM ARMY RESEARCH LABORATORY

Monitoring Russian Online Information Operations: An OssaLabs Case Study

by Timur Chabuk, Adam Jonas, and Sue Kase

Approved for public release; distribution is unlimited.

NOTICES

Disclaimers

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof.

Destroy this report when it is no longer needed. Do not return it to the originator.

ARL-TR-8732 ● JULY 2019

Monitoring Russian Online Information Operations: An OssaLabs Case Study

by Timur Chabuk
Intelligent Information Systems, Perceptronics Solutions Inc

Adam Jonas
Network Engagement Team, Army Training and Doctrine Command (TRADOC)

Sue Kase
Computational and Information Sciences Directorate,
CCDC Army Research Laboratory

Approved for public release; distribution is unlimited.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.
PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED (From - To)
July 2019	Technical Report	January 2017–May 2019

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
Monitoring Russian Online Information Operations: An OssaLabs Case Study	W911QX17C0007
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)	5d. PROJECT NUMBER
Timur Chabuk, Adam Jonas, and Sue Kase	5e. TASK NUMBER
	5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	8. PERFORMING ORGANIZATION REPORT NUMBER
CCDC Army Research Laboratory ATTN: FCDD-RLC-NC Aberdeen Proving Ground, MD 21005	ARL-TR-8732

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES

ORCID ID: Timur Chabuk, 0000-0002-5793-4678

14. ABSTRACT

Full-spectrum information operations (IO) involve coordinated action across a diverse set of communication media ranging from cyber-attacks, to press releases, to sponsored advertisements in magazines. Social media is one particularly important medium that offers unprecedented opportunity for IO to reach large-scale numbers of people. Russian IO are using social media to great success against the US and its allies. These operations have created public confusion and discord domestically while preparing the battlefield for more traditional maneuvers in US Army areas of interest. This report introduces a new technology for monitoring social media for IO. The OssaLabs platform allows analysts to scan large swathes of social media space to discover key narratives, identify important actors, and track ongoing IO. The major features of the OssaLabs social media monitoring platform are described. Then, a case study is presented in which the OssaLabs platform was used to monitor social media space in the Baltic region, where Russia is believed to be conducting intensive IO. We conclude with a discussion of key findings from the case study and identify important directions for future technology development.

15. SUBJECT TERMS

information operations, social media monitoring, narratives, Russian language, Baltic region

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Sue Kase
Unclassified	Unclassified	Unclassified	UU	23	19b. TELEPHONE NUMBER (Include area code) 410-278-7009

Standard Form 298 (Rev. 8/98)
Prescribed by ANSI Std. Z39.18

Contents

List of Figures

1. Introduction

State and nonstate actors around the world are increasingly carrying out information operations (IO) to shape public opinion and influence world events. Full-spectrum IO involve coordinated action across a diverse set of different media, ranging from cyber-attacks, to press releases, to simple, sponsored advertisements in magazines. Social media is one particularly important medium in which IO are conducted.

Modern social media platforms offer unprecedented opportunities for IO to reach large-scale numbers of people with messaging that can be micro-targeted to individual users and groups. Recent events, including Russian interference in the United States (US) Presidential election (Bessi and Ferrara 2016), Russian IO conducted in support of military operations in Ukraine (Jaitner 2015), Islamic State in Iraq and Syria (ISIS) recruitment and radicalization through social media, and Iranian influence operations targeting the American public, have drawn public attention to this issue. Furthermore, recent studies suggest that numerous countries around the world actively participate in manipulating social media space (Bradshaw and Howard 2017). Therefore, it is essential the US and its allies develop technologies to detect, analyze, and monitor social media IO (Chabuk and Jonas 2018).

This report introduces a new technology for monitoring social media space for IO. The OssaLabs platform allows analysts to easily monitor large swathes of social media space to discover key narratives, identify important actors, and monitor ongoing IO. In the remainder of this report, we discuss the challenges involved with monitoring social-media-based IO with a particular focus on Russian IO. We then describe the OssaLabs social media monitoring platform and its major features. Next, we discuss a case study in which the OssaLabs platform was used to monitor social media space in the Baltic region where Russia is believed to be conducting intensive IO (Radin 2017; Standish 2017). The final section summarizes key findings from this case study and identifies important directions for future technology development.

2. Challenges of Monitoring IO

The social media information environment is challenging to understand because of its overwhelming complexity and constant evolution. Social media space is shaped by a wide range of actors, each with their own objectives, strategies, and tactics. Narratives and counter-narratives are created, amplified, and countered continuously in rapid response to unfolding events. Various groups and communities respond to narratives in different ways, further shaping the

information environment through their likes, shares, and comments. Effectively monitoring social media space to detect and understand ongoing IO requires capabilities to identify all the previously stated dynamic elements in real time. In the remainder of this section, we discuss some of the objectives, strategies, tactics, and actors that are relevant to Russian social media IO.

Russia carries out IO with two main objectives: 1) to weaken Western nations and Western-led international organizations such as the United Nations and North Atlantic Treaty Organization (NATO); and 2) to increase Russia's international influence, particularly in neighboring former Soviet countries (Darczewska 2014).

To weaken other countries, Russia employs strategies that seek to exploit and exacerbate existing divisions within the country. For example, Russia used paid Facebook advertising to promote both sides of divisive issues, such as promoting the Black Lives Matter social movement while also paying to promote advertisements that labeled them a dangerous threat (Byers 2017). Russian IO also seeks to undermine popular support for Western-led institutions such as the European Union (e.g., Brexit) and NATO. Target governments are often accused of being incompetent or corrupt, and in some cases (e.g., Lithuania) Russian IO question the historical legitimacy of the very existence of countries. Collectively, these strategies undermine the cohesiveness of the target nation's population, thereby making it more difficult for the target nation to respond to Russian provocations. As others have noted, the ultimate accomplishment would be to create enough confusion and dissent so as to prevent a successful Article 5 vote among NATO nations (Giles 2016). Russia also uses IO in support of kinetic operations (e.g., Ukraine), to create enough confusion and misinformation so that an adversary's decision-making is altered or delayed.

In addition to targeting NATO and Western-allied nations, Russia targets former Soviet republics with large Russian-language speaking populations in an effort to reassert Russian influence by turning the Russian diaspora against the target government and strengthening their allegiance to Russia (Jaitner and Mattsson 2015). To accomplish this, Russia uses a mix of strategies, including soft strategies such as promoting nostalgia for the Soviet era by reminiscing about Soviet-era actresses, and hard strategies such as spreading fake news stories about pending legislation that will allegedly ban the Russian language.

Other Russian influence tactics shape the social media information environment through a variety of media-generated narratives. News stories containing desired narratives are created by media outlets that are effectively controlled by the Russian government (e.g., RT [formerly Russia Today] and Sputnik). Covert intelligence agencies (e.g., the Internet Research Agency) employ full-time operatives to create

content in the form of graphics, videos, memes, and fake news stories promoting desired narratives that are designed to resonate with target civilian populations (Bradshaw and Howard 2017). Fake news stories in particular have spread further and faster than real news stories (Vosoughi et al. 2018) and are sometimes repeated by mainstream news media and political candidates as if they are true (Lazar 2017). Covert operatives also operate bots, sockpuppets, and so-called "troll armies" to amplify any desired narratives through sharing and liking of posts. Information operations are not only carried out by covert operatives. Official government social media accounts such as foreign ministries, embassies, and diplomats often lend the appearance of legitimacy to fake news stories by sharing and liking them. The ultimate goal is for these narratives to gain traction with target civilian communities who will then further spread the narrative and influence public opinion.

In addition to amplifying desired narratives, Russia spends considerable time and resources on counter-narratives and information contrary to Russian interests. One widely recognized way of understanding Russian disinformation operations is by categorizing them into four classes of tactics: *dismiss*, *distort*, *dismay*, and *distract* (Nimmo 2015). Dismiss tactics are used to undermine belief in facts that are contrary to Russian interests by simply denying their truth or by attacking the credibility of the messenger even in the face of clear supporting evidence. If dismiss tactics are not effective, then distort tactics are used to modify information by cherry-picking facts and adding lies to otherwise true information. Dismiss and distort tactics can be effective even when it is plainly evident that the facts being dismissed are true and the distortions are false. All that is required of these tactics is to "muddy the waters" enough to sow doubts in the target population and undermine support for the target government (Jaitner and Mattsson 2015).

Dismay tactics are used to simply intimidate and scare the public by activating emotionally charged fears and anxieties through overblown threats and rhetoric, and to silence critics by intimidating them with harassments and personal threats. When dismiss, distort, and dismay tactics fail, distraction techniques are used to change the subject or topic away from the unfavorable information. This is often accomplished by promoting wildly sensational fake news stories and attention-grabbing headlines. Strategically, this can take the form of "hashtag poisoning" in which a hashtag that is being used by activists to influence opinion and coordinate action is made unusable by flooding social media with nonsense posts containing that same hashtag (Bradshaw and Howard 2017).

3. OssaLabs Social Media Monitoring Platform

Considering the many types of strategies and influencing tactics, the end goals of Russian online IO are often difficult to distill, especially in the short term. In the past, Russia demonstrated how a relatively small, but well-coordinated, investment of capital can have disproportionate effects on adversaries. It is therefore imperative that the US government fund the development of new technologies to identify and mitigate such threats. This section describes the OssaLabs social media monitoring platform and its capabilities for detecting and tracking IO.

3.1 Overview

The OssaLabs platform was designed specifically to help analysts identify, analyze, and monitor social media IO. OssaLabs was developed through funding by the Small Business Innovation Research Program from the Defense Advanced Research Projects Agency, Office of Naval Research, US Army Combat Capabilities Development Command Army Research Laboratory (CCDC ARL), and other agencies tasked with protecting US national security. OssaLabs is currently collaborating with the US Army Training and Doctrine Command Intelligence Directorate (TRADOC G2) to transition the platform to the Army at large with additional funding assistance from CCDC ARL.

The OssaLabs platform is a cloud-hosted web application that runs seamlessly in modern browsers across major computer operating systems. OssaLabs enables a full social media assessment workflow, including collection of social media data, organization and analysis of the data, and visualization of the analysis results to support user understanding. The three major components of the OssaLabs platform are *Queries*, *Groups*, and *Dashboards*. Analysts create queries to collect data of interest from various social media platforms. Groups are used to organize and index the collected social media data as customized by the analyst. Dashboards provide analysts with a polished presentation of the collected data and analysis results. Each is described in more detail as follows.

3.2 Queries

Analysts create and configure *queries* to collect data of interest from social media platforms. Currently, there are two types of queries for collecting data from Twitter.

1) **Twitter Keyword Queries** allow analysts to collect tweets that contain specific keywords or keyword combinations. Keywords can be individual words, phrases, hashtags, user handles, or URLs. Analysts specify a set of "include" keywords, and any tweet that contains at least one of those

keywords will be collected. Additionally, analysts can optionally further refine their query by specifying "require" and "exclude" keywords. If any "require" keywords are entered, then a tweet must contain at least one "include" and one "require" keyword in order to be collected. If any "exclude" keywords are specified, then any tweets that contain any "exclude" words will not be collected.

2) **Twitter User Queries** allow analysts to collect tweets from specific Twitter user accounts. All tweets from the specified accounts will be collected, regardless of keywords used. This is a useful feature for tracking known and important user social media accounts.

We are in the process of developing new query capabilities for collecting data from other social media platforms such as Facebook and VKontakte, a Russian online social media and networking service.

Queries in foreign languages are seamlessly supported by the OssaLabs platform. To collect foreign language social media data, the analyst simply selects the desired language from a drop-down list and enters the desired query parameters (e.g., keywords) in that language. The OssaLabs platform will then collect matching tweets in that language and automatically translate them into English. To perform the translations, OssaLabs is currently integrated with Microsoft Azure translation services.

3.3 Groups

Analysts can organize and index collected data by creating and configuring *groups*. There are currently three distinct types of groups that are supported by the OssaLabs platform:

1) **Mention Groups** allow analysts to specify a list of words and phrases that capture some concept or meaning. For example, an analyst can create a mention group named "Coffee" that can include the phrases "java," "cup of joe," and "coffee".

2) **Participant Groups** allow analysts to define a list of Twitter accounts of interest. For example, an analyst may maintain a participant group for known troll accounts, or for journalists who write about a particular topic. Once defined, participant groups are useful for sorting, organizing, and filtering data.

3) **Follower Groups** allow analysts to create a set of Twitter accounts that follow one or more of a specified set of Twitter accounts. For example, an

analyst may be interested in the followers of a particular foreign leader. The analyst can create a follower group by specifying a foreign leader's Twitter handle. The OssaLabs platform will obtain from Twitter the list of accounts that follow that handle and store them as members of a follower group. The collected tweets can later be filtered, sorted, or organized according to this follower group.

3.4 Dashboards

Analysts can review collected data and analysis results in *dashboard*s. To create a new dashboard, the analyst specifies what data are to be collected by selecting specific queries or mention groups, or by specifying specific phrases or keywords to include. More targeted analyses can be performed by creating complex criteria specifying grammar expressions that reference additional qualities of tweets (e.g., retweet or not) and analysis results (e.g., sentiment of tweets).

Analysts specify the date range of collected data they wish to view by using a calendar drop-down list. In the case of foreign language data, the analyst can also specify whether they want to view the dashboard in the original language or translated to English. Filtering controls within the dashboard allow the analyst to quickly explore various options for further refining the data, including filtering by keywords, sentiment, mention groups, participant groups, and follower groups. An example OssaLabs dashboard is shown in Fig. 1.

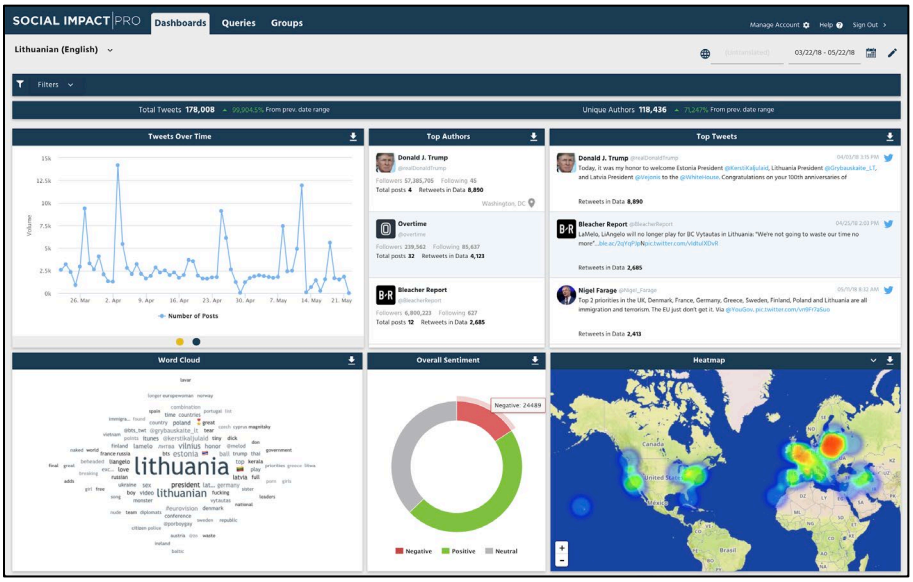

Fig. 1 Analysts use OssaLabs dashboards to review, filter, and visualize collected data and analysis results

Dashboards offer a suite of analysis and visualization capabilities to assist the analyst in understanding the social media space. Figure 1 shows several of the subwindows contained in a dashboard: Tweets Over Time, Top Authors, Top Tweets, Word Cloud, Overall Sentiment, and Geospatial Analysis. A brief description of each subwindow follows:

1) *Tweets Over Time* subwindow (Fig. 1, upper left) displays a timeline showing the volume of tweets. This capability allows analysts to easily identify trends over a selected time period as well as anomalies such as lulls and spikes in activity.

2) *Top Authors and Top Tweets* subwindows (Fig. 1, upper center and right) display a list of three authors whose tweets were retweeted the most, and a list of three tweets that were retweeted the most, respectively. These capabilities provide the analyst with a quick at-a-glance understanding of what is occurring right now in social media.

3) *Word Cloud* subwindow (Fig. 1, lower left) displays the most commonly used words in the tweets collected so far. This graphic is particularly useful in discovering words that are commonly used alongside the analyst-defined keywords.

4) *Overall Sentiment* subwindow (Fig. 1, lower center) automatically analyzes the sentiment of all collected tweets and visualizes the total number of positive, neutral, and negative tweets in the form of a donut chart. The overall sentiment graphic provides a top-level understanding of the emotional aspects of the tweet content. OssaLabs implements the overall sentiment analysis by using the VADER sentiment tool (Hutto and Gilbert 2014).

5) *Geospatial Analysis* subwindow (Fig. 1, lower right) identifies the location of the tweets based on any geotag data available on the tweet, as well as geographic information from the author's Twitter profile. The results of the analysis are visualized using a heat map, allowing analysts to quickly understand the geographic distribution of ongoing conversations.

6) *Topic Detection* subwindow (not shown in Fig. 1) analyzes all collected tweets and automatically clusters them into topics based on their semantic similarity using algorithms inspired by KeyGraph (Sayyadi et al. 2009). The visualization format presents the identified tweet clusters grouped together based on the similarities between the clusters. This capability offers analysts information on how many tweets are in each cluster, and a representative tweet example from each cluster that appears on top. These clusters allow

analysts to quickly scan large amounts of data without having to examine each tweet individually. Additional clustering metrics are available such as number of tweets, number of Twitter users, and tweet sentiment.

4. Case Study: Using OssaLabs to Monitor Baltic Social Media Space

In this section, we present a case study where the capabilities provided by the OssaLabs platform are used to monitor the Baltic social media space. There is a brief background section explaining the ongoing tensions between Russia and the Baltic nations. The second section describes a step-by-step process for using OssaLabs to monitor this social media space for Russian IO.

4.1 Scenario Background

The Baltic region is a hotbed of tension between Russia and the West. Historically, the countries of Latvia, Estonia, and Lithuania were all part of the Soviet Union. After the fall of the Soviet Union, these countries left and eventually became members of NATO. Russia wants to reassert influence in these countries and undermine their membership in NATO. Russia's efforts toward this strategy include a wide range of kinetic, economic, cyber, and IO. It is widely believed that Russia promotes various narratives in social media space that are designed to reinforce Russian-language-speaking people's identification with Russia and to undermine Baltic nation citizens' confidence in their own country. These narratives include promoting nostalgia from the Soviet era, questioning and undermining the historical legitimacy of the Baltic nations, exacerbating ethnic tensions between Russian-language speakers and others, and alleging incompetence and corruption of Baltic nation governments and leaders.

4.2 Monitoring Walkthrough

We begin our case study as an analyst assigned the mission of monitoring the Baltic social media space for Russian IO activity. We first create a Twitter keyword query targeting high-level terms that are relevant to Lithuania, including the English words "Lithuania," "Lithuanian," "Klaipeda," and "Vilnius". Then, we create three additional Twitter keyword queries to collect these same terms, but in the Russian, Lithuanian, and Polish languages. Figure 2 shows a Twitter keyword query using Russian-language keywords related to Lithuania. Similar queries could be created for Estonia and Latvia as well, although in this case study we only focus on Lithuania. Data were collected using these keyword queries starting in mid-March 2018 and are continuing to be collected today.

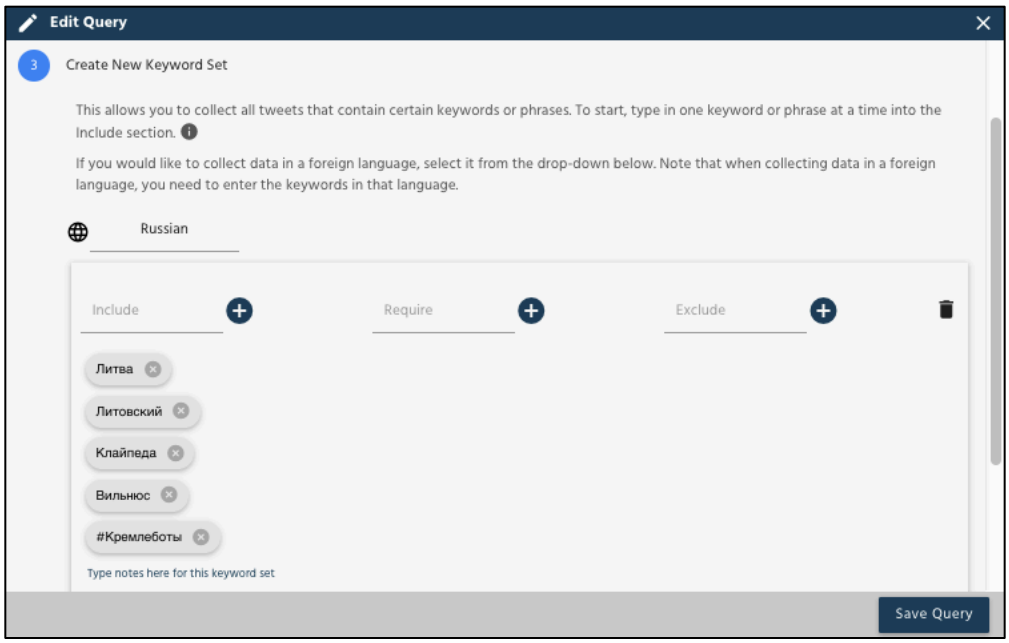

Fig. 2 Twitter keyword query configured to collect Russian-language tweets containing various terms related to Lithuania

Because we are primarily interested in Russia's suspected targeting of Russian-language-speaking populations in the Baltics, we focus our attention on the Russian-language data collected from the keyword queries. Next, we create a dashboard that pulls in all data collected from the queries. We name this dashboard the "Lithuania (Russia)" dashboard. Once in the "Lithuania (Russia)" dashboard, we select a two-month date range of 22 March–22 May 2018.

The analysis process begins with reviewing the top tweets and authors within the date range. This immediately reveals some interesting content that seems relevant to pro-Kremlin narratives and counter-narratives that we would expect to be playing out in the social media space. For example, Fig. 3 shows the top three retweeted tweets in the date range. A brief commentary of potential IO narratives follows:

1) In the @Vitauskas_A tweet (Fig. 3, left side), the author warns that had Lithuania not joined NATO, Russian troops would have invaded Lithuania under the pretense of protecting Russian-language-speaking peoples.

2) At first glance, the @Current_Policy tweet (Fig. 3, center) appears to be expressing support or admiration for Lithuanian Special Forces. However, if we also view the video contained in the tweet, it becomes clear that the text is actually sarcastically mocking the poor performance of the Soldiers attempting over and over again to knock down a door.

3) In the @ClownIT tweet (Fig. 3, right side), we are unsure how to interpret the message and do not understand why this tweet is so widely retweeted. We suspect that some kind of joke is at play that is lost in the translation process.

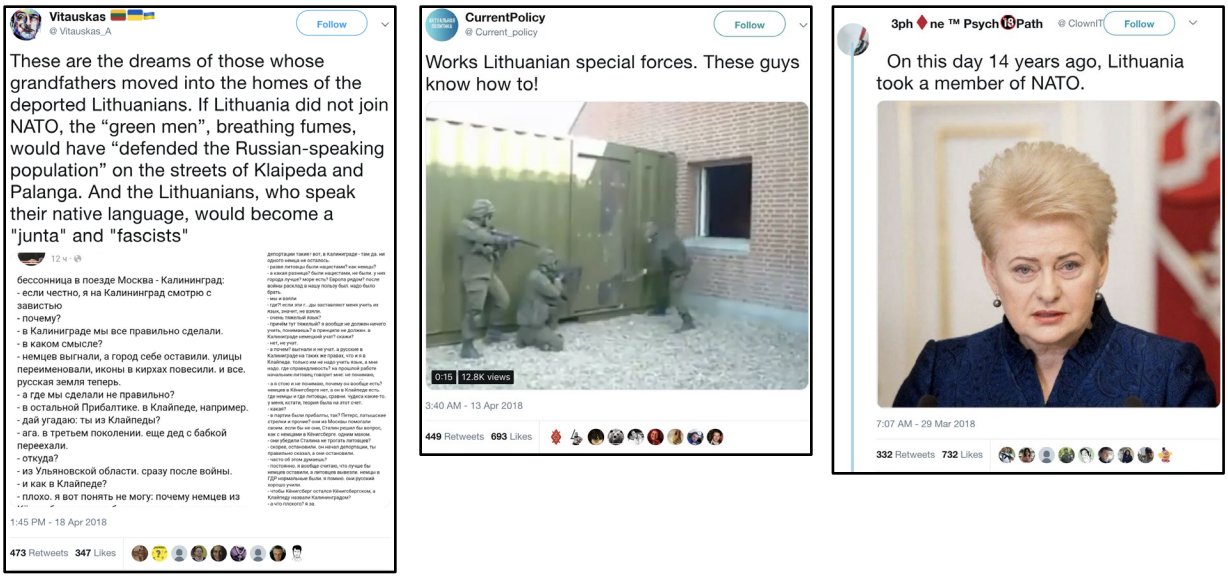

Fig. 3 Three most retweeted tweets collected during the date range

To further refine the suspected narratives gleaned from analyzing the top tweets and authors, we use the OssaLabs topic detection with clustering capabilities. Two noteworthy topics are discovered and summarized here:

1) The *Forest Brothers* topic cluster (Fig. 4, left side) is exemplified by another tweet from the @Vitauskus_A Twitter handle. In this tweet, a collage of war-torn Lithuania during World War II is shown. The pictures are described as 1) showing damage done by the Soviet Union when they invaded Lithuania to simultaneously expel the Germans and occupy the country; 2) dead members of the militia group known as the “Forest Brothers,” who were young men of the Baltics who fought against the Soviet occupation. The Forest Brothers emerged as a great source of national pride for Lithuanians; and 3) mass, forced deportations of Lithuanian people to Siberia. Collectively, this topic cluster appears designed to convey the message that Russia is not a historical friend to Lithuania.

2) The *Return the Vilnius* topic cluster (Fig. 4, right side), on the other hand, shows a group of tweets exemplified by the pictures tweeted by Nikolai Starikov featuring a lengthy video arguing that a major Lithuanian city,

Vilnius, should be returned to Russia. This topic of tweets appears designed to undermine Lithuanian sovereignty by promoting the idea that half of Lithuania "rightfully" belongs to Russia.

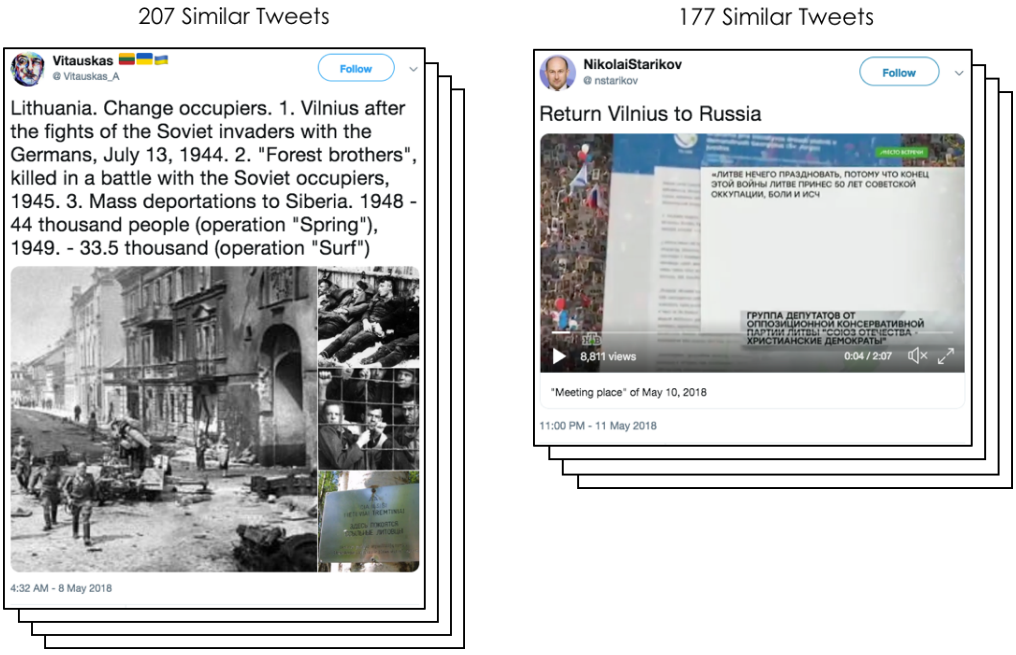

Fig. 4 Representative tweets from the two largest topic clusters identified in the collected data

The combined capabilities on the OssaLabs platform appear to point to three accounts that are clearly pro-Kremlin. These accounts routinely push talking points and narratives promoting a Kremlin point of view. We also identified three anti-Kremlin accounts that actively counter those narratives and promote criticism of Russia. The identification of these six accounts is confirmed by social network analysis (SNA).

SNA can quantitatively help answer questions such as who has direct influence online; who are the authors influencing the influencers; what subgroups within a broader network represent social/ideological divides; and how many members of the network are disproportionately targeted or elevated by bot activity (Jonas 2017). The OssaLabs platform provides SNA capabilities that examine patterns of interactions between social media users to discover communities and influencers. Specifically, the OssaLabs platform is used to construct a social network with these characteristics: 1) each node in the network is a social media user who appeared in the data; 2) social media users are connected by an edge if they retweeted the same tweet; and 3) the edge weight reflects how many times such co-retweeting occurs. This social network is exported from OssaLabs and visualized using Gephi

(Bastian et al. 2009), as shown in Fig. 5. Community detection algorithms (Blondel et al. 2008) are applied to identify distinct parts of the network that are more tightly connected to each other, forming a *community*. Each distinct community is randomly assigned a color. OssaLabs is used to manually review the accounts and tweets in each community, resulting in the labeling of two important communities. The red-colored community is pro-Kremlin and the blue-colored community is anti-Kremlin. Examination of the three most retweeted users within each community group confirms the same accounts that we had previously identified using the top tweets and authors and the topic detection OssaLabs capabilities.

We can use the follower group capability for creating anti-Kremlin and pro-Kremlin groups of users, which we can then apply a filter to in the dashboard. For example, by selecting the anti-Kremlin follower group from the filter drop-down list, we can view only content that is from social media users who follow the accounts we have identified as being anti-Kremlin. And we can apply a similar filtering process for the pro-Kremlin accounts. This filtering capability assists analysts in maintaining an awareness of the trending topics and narratives within each of these groups on a daily basis.

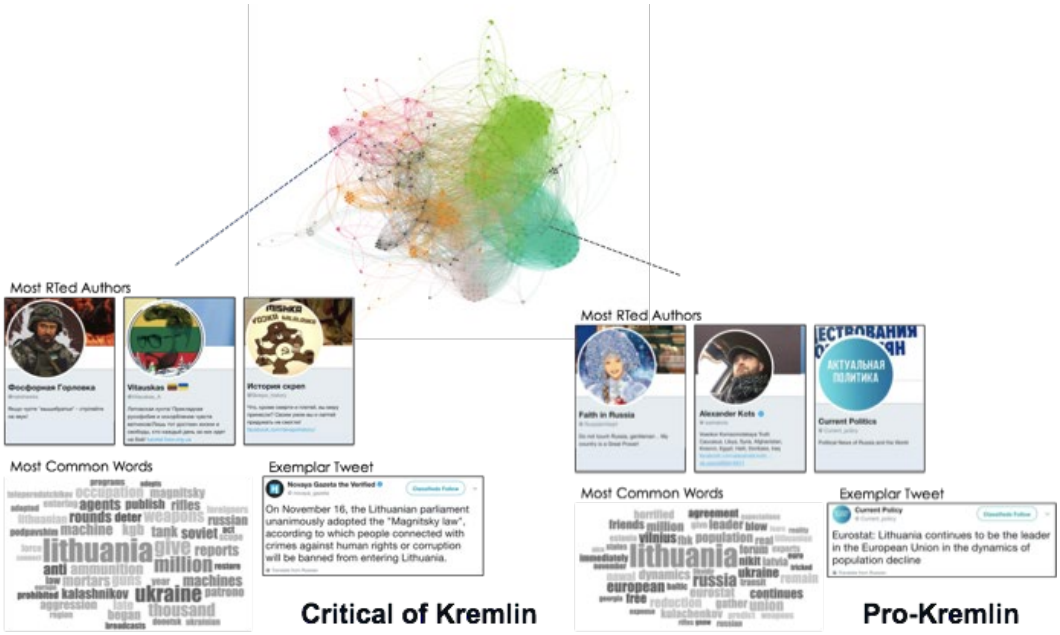

Fig. 5 Social network analysis reveals two distinct groups among several different groups: one group is pro-Kremlin (red colored) and the other group is anti-Kremlin (blue colored)

5. Conclusion

Information operations, including social-media-based IO, will continue to be an increasingly important part of modern warfare. It is essential that the US and its allies continue to develop new technologies to identify, analyze, and counter social media IO. This report describes and demonstrates a new technology called OssaLabs that can be used to monitor large swathes of social media space, discover key narratives, identify key actors, and use that information to further focus ongoing monitoring efforts.

There are several important directions for future development of OssaLabs and related platforms. First, enhanced support for the detection and characterization of communities is needed. As described in the case study section, narratives are designed to resonate with targeted communities. It is only by understanding those communities and what defines them (i.e., their shared beliefs and perspectives) that we can successfully identify and monitor any and all IO. A second important extension is the development of automated support for the detection of important actors. This includes influential accounts within civilian communities, operative-controlled accounts that are seeding and promoting narratives, troll accounts that are amplifying messaging and harassing critics, and even official government accounts that are promoting and lending legitimacy to misinformation. Lastly, the development of customized analyses and dashboards to answer specific questions is an important next step to easily extracting value from social media analytics. For example, a suite of analyses and dashboards can identify when major critics of Russian policy are being harassed online and who the harassers are.

Russian IO against the US is not new by any means. Russia quickly evolves their use of IO by leveraging social media and will continue to use social-media-based tactics to maintain domestic power while creating political discord among adversaries and supporting traditional military maneuvers in locations such as the Ukraine. Without mastery of new online detection technologies that enable broader thinking about military maneuvers, our IO overmatch cannot be maintained, which will result in Russia and other countries employing similar IO tactics and gaining technological and military superiority in the very near future.

6. References

Bastian M, Heymann S, Jacomy M. Gephi: an open source software for exploring and manipulating networks. In: Proceedings of the Third International AAAI Conference on Weblogs and Social Media; 2009 May 17–20; San Jose, CA. Palo Alto (CA): Association for the Advancement of Artificial Intelligence; c2009. p. 361–362.

Bessi A, Ferrara E. Social bots distort the 2016 US Presidential election online discussion. Chicago (IL): University of Illinois at Chicago. *First Monday*; 2016;21(11) [accessed 2019 May 30]. http://firstmonday.org/article/view/7090/5653.

Blondel VD, Guillaume JL, Lambiotte R, Lefebvre E. Fast unfolding of communities in large networks. J Stat Mech-Theory. 2008;10:10008.

Bradshaw S, Howard PN. Troops, trolls and troublemakers: a global inventory of organized social media manipulation. Oxford (England): University of Oxford; c2017. Computational Propaganda Research Project; Working Paper No. 2017.12.

Byers D. Exclusive: Russian-bought Black Lives Matter ad on Facebook targeted Baltimore and Ferguson. CNN Business; 2017 Sep 28 [accessed 2019 May 30]. http://money.cnn.com/2017/09/27/media/facebook-black-lives-matter-targeting/index.html.

Chabuk T, Jonas AB. Understanding Russian information operations. SIGNAL. 2018 Sep 1 [accessed 2019 May 30]. https://www.afcea.org/content/understanding-russian-information-operations.

Darczewska J. The anatomy of Russian information warfare. The Crimean operation, a case study. Point of View. Warsaw (Poland): Centre for Eastern Studies. 2014;(42):1–37. http://www.osw.waw.pl/sites/default/files/the_anatomy_of_russian_information_warfare.pdf.

Giles K. Handbook of Russian information warfare. NATO Defense College Fellowship Monograph. Rome (Italy): NATO Defense College, Research Division; 2016 Nov [accessed 2019 May 30]. http://www.ndc.nato.int/news/news.php?icode=995.

Hutto CJ, Gilbert E. VADER: a parsimonious rule-based model for sentiment analysis of social media text. In: Proceedings of the Eighth International AAAI Conference on Weblogs and Social Media (ICWSM-14); 2014 June 1–4; Ann Arbor (MI). Palo Alto (CA): Association for the Advancement of Artificial

Intelligence; c2014. p. 216–225. [accessed 2019 May 30]. http://comp.social.gatech.edu/papers/icwsm14.vader.hutto.pdf.2014.

Jaitner M. Russian information warfare: lessons from Ukraine. In: Geers K, editor. Chapter 10, Cyber War in Perspective: Russian Aggression against Ukraine. Tallinn (Estonia): NATO CCD COE Publications. c2015. 87–94.

Jaitner M, Mattsson PA. Russian Information warfare of 2014. In: Maybaum M, Osula A-M, Lindström L, editors. CyCon 2015. Proceedings of the 2015 7th International Conference on Cyber Conflict: Architectures in Cyberspace; 2015 May 26–29; Tallinn (Estonia): NATO CCD COE Publications. c2015. p. 39–52. IEEE.

Jonas AB. How the hashtag is changing warfare: armies of social media bots battle for hearts and minds online. SIGNAL. 2017 June 29 [accessed 2019 May 30]. https://www.afcea.org/content/how-hashtag-changing-warfare.

Lazar D. Keynote presentation at NASN2017: 1st North American Social Networks Association (NASN) Conference of the International Network for Social Network Analysis; 2017 July 23–30; Washington, DC.

Nimmo B. Anatomy of an info-war: how Russia's propaganda machine works, and how to counter it. Bratislava, Slovakia: Central European Policy Institute; 2015 May 19 [accessed 2019 May 30]. https://www.stopfake.org/en/anatomy-of-an-info-war-how-russia-s-propaganda-machine-works-and-how-to-counter-it.

Radin A. Hybrid warfare in the Baltics: threats and potential responses. No. RR-1577-AF. Santa Monica, CA: RAND Corporation; 2017 [accessed 2019 May 30]. https://www.rand.org/pubs/research_reports/RR1577.html.

Sayyadi H, Hurst M, Maykov A. Event detection and tracking in social streams. In: Proceedings of the Third International AAAI Conference on Weblogs and Social Media; 2009 May 17–20; San Jose, CA. Palo Alto (CA): Association for the Advancement of Artificial Intelligence; c2009. p. 311–314.

Standish R. Russia's neighbors respond to Putin's 'Hybrid War'. Washington (DC): Foreign Policy. 2017 Oct 12 [accessed 2019 May 30]. https://foreignpolicy.com/2017/10/12/russias-neighbors-respond-to-putins-hybrid-warlatvia-estonia-lithuania-finland.

Vosoughi S, Roy D, Aral S. The spread of true and false news online. Science. 2018;359(6380):1146–1151.

List of Symbols, Abbreviations, and Acronyms

CCDC ARL	US Army Combat Capabilities Development Command Army Research Laboratory
IO	information operations
ISIS	Islamic State in Iraq and Syria
NATO	North Atlantic Treaty Organization
SNA	social network analysis
TRADOC G2	US Army Training and Doctrine Command Intelligence Directorate
US	United States
URL	Uniform Resource Locator

1 (PDF)	DEFENSE TECHNICAL INFORMATION CTR DTIC OCA
2 (PDF)	CCDC ARL IMAL HRA RECORDS MGMT FCDD RLD CL TECH LIB
1 (PDF)	GOVT PRINTG OFC A MALHOTRA
1 (PDF)	CCDC ARL FCDD RLC NC S KASE