EXAMINING STRATEGIC INTEGRATION OF SOCIAL MEDIA PLATFORMS IN DISINFORMATION CAMPAIGN COORDINATION

Nitin Agarwal, Kiran Kumar Bandeli

Abstract

Social media platforms are widely used for sharing information. Although social media use is generally benign, such platforms can also be used for a variety of malicious activities, including the dissemination of propaganda, hoaxes, and fake news to influence the public. The availability of inexpensive and ubiquitous mass communication tools has made such malicious activity much more convenient and effective. In this paper we study how blogs act as virtual spaces where malicious narratives are framed and then further disseminated through social media platforms such as Twitter and Facebook. To discover how such disinformation campaigns work, it is necessary to examine the link between blogs and social media platforms and the role they play in media orchestration strategies, more specifically cross-media and mix-media strategies. We have carried out an in-depth examination of information networks, using social network analysis and cyber forensics, to identify prominent information actors and the leading coordinators of several disinformation campaigns. The research methodology we have developed reveals a massive disinformation campaign pertaining to the Baltic region, conducted primarily through blogs but strategically linking to a variety of other social media platforms, e.g. Twitter, Facebook, YouTube, and VKontakte.

Note: Although blogs fall under the broad definition of social media, for the purposes of this paper we distinguish between blogs and social media, using the term social media to refer to forums that have some sort of formal membership, but where any member can post about anything (within certain limits of legality and decency) to the audience of the account holder's choice; a blog is the product of an individual (at least ostensibly), writing about a set of topics limited by the blogger's interests.

Keywords—*social media, blogs, social network analysis, cyber forensics, disinformation, mix-media, cross-media.*

About the Authors

Nitin Agarwal is the Jerry L. Maulden-Entergy Endowed Chair and Distinguished Professor of Information Science at the University of Arkansas and Director of the Collaboratorium for Social Media and Online Behavioral Studies (COSMOS). He researches social computing, cyber information campaigns, deviant behavior modeling, group dynamics, social-cyber forensics, data mining, and privacy.

Kiran Kumar Bandeli is a PhD student at the University of Arkansas at Little Rock. He majors in Computer and Information Sciences with research interest in social media analysis, social computing, narrative analysis, and mis/disinformation campaigns.

Introduction

The power of blogs and social media as positive vehicles of transformation first became apparent around 2011, during the many events of the Arab Spring and subsequent Occupy demonstrations worldwide. However, this power has been recently harnessed by state and non-state actors, extremists, and terrorist groups to influence online discourse, steer mass thinking, and polarise communities, posing a dangerous force against democracy. This power, coupled with the proliferation of high-speed Internet-enabled mobile devices, has propelled us into an age where a lie can travel around the world several times to reach our screens before truth-seekers have had a chance to verify the claim.[1]

1 Soroush Vosoughi, Deb Roy, and Sinan Aral, 'The Spread of True and False News Online', *Science* 359, № 6380 (2018): 1146–51.

Weblogs, or blogs, are first-person diary-entry-style postings presented in reverse chronological order. Because they cost nothing to set up, blogs have become one of the most popular platforms people use to voice their opinions. Blogs have become the preferred mode of carriage for all other media— everything digital (photos, videos, etc.) migrates to blogs. Consumers of information have now also become information producers and distributors. This practice has led to a meteoric rise in the popularity of blogs, and in their influence on political communication. Blogs are among the fastest-growing web-based forums.[2] At the start of 2003 there were fewer than 1 m blogs in existence; over the next three years their number doubled every six months. By 2011, there were 158 m blogs, with more than 1 m posts being produced every day. According to recent estimates, this number has grown to over 3 m new posts added daily to the blogosphere.[3] According to WordPress, over 409 m people view close to 22 bn pages each month.[4] Most assume that blogs empower ordinary citizens and expand the social and ideological diversity of the voices online. Blogs are thought to make political discourse less exclusive, making them fertile ground for ideal citizen journalism. These beliefs are mistaken.[5]

Blogs are free to use, barriers to publication are low, and the information environment of the blogosphere is casual and does not require scrupulous fact-checking. These factors put the quality of the information presented at great risk. Although there are several million blogs, only a tiny fraction of them maintain a large readership. Moreover, bloggers often do not represent the opinions of the broader electorate. Blogs have done little to amplify the political voice of average citizens; instead, certain partisan narratives are promoted. The low barrier to publication guarantees that anyone with something to say—be it true or false, fact or opinion, fair or biased—can post it to a blog for the world to see. Today, when it is more convenient for people to get their news from blogs and social media rather than the mainstream media, irresponsible citizen journalism poses a threat to democratic principles and institutions by misrepresenting facts and information. The Pew Research Center conducted a poll in the US regarding reputable sources of information and reported that over one-third (34%) of respondents trust news from social media sources,[6] and over three-fourths (77%) trust information from friends and family, shared through a variety of means including blogs and social

2 Kenneth Jost and Melissa J. Hipolit, *Blog Explosion: Are Blogs a Passing Fad or a Lasting Revolution?*, *CQ Researcher* 16 (2006), Issue 22 (9 June). PAYWALL.
3 Blogging statistics: blog posts written today, is a real-time counter showing the global activity on wordpress bloggers. [last accessed 4 June 2018].
4 Stats: A Look at Activity Across WordPress.com, provides a visualisation of current posts, comments, and likes. [last accessed 4 June 2018]
5 Matthew Hindman, *The Myth of Digital Democracy* (Princeton, NJ: Princeton University Press, 2008).
6 Michael Barthel and Amy Mitchell, 'Americans' Attitudes About the News Media Deeply Divided Along Partisan Lines', *Pew Research Center's Journalism Project* (blog), 10 May 2017.

media channels. These statistics show the dangerous power of bloggers and social media users to conduct influence operations that manipulate public discourse.

Plenty of empirical evidence exists demonstrating the role blogs have played in the constant barrage of fake news and misinformation during various geopolitical events, regional as well as global, over the last several years. One of the most prominent examples of this phenomenon is the 2016 US Presidential election,[7] when a number of misinformation-riddled stories were planted in clickbait[8]-driven post-truth[9] media for financial and political incentives. Macedonian teens tapped into the digital gold rush by setting up blogs with content copied verbatim from alt-right[10] news sites. During the Ukraine-Russia crisis, sites like LiveJournal, various blogging platforms, and VKontakte (a Russian social media platform), were used as propaganda machines justifying the Kremlin's policies and actions.[11] According to *Interpret Magazine*, the Kremlin recruited over 250 trolls,[12] each of whom was paid $917 per month to work round the clock producing posts on social media and comments on mainstream media outlets. These trolls managed a stream of invective against pro-Ukrainian media and Western news stories containing unflattering information about Russia, using multiple 'sock puppet' accounts—online identities used for deception. These troll armies, or 'web brigades', piggyback on the popularity of social media to disseminate fake pictures and videos promoting their ideological goals. They have coordinated some very effective disinformation campaigns, to which even legitimate news organisations fall prey. To stem the tide of information subversion, or at least to raise awareness, several online crowdsourcing-based efforts, such as StopFake.org and euvsdisinfo.eu, have been created to identify and debunk deliberately misleading imagery and stories about the war in Ukraine. However, such efforts are severely limited and the fact checkers are easily outnumbered by the vast troll armies.

7 Samanth Subramanian, 'The Macedonian Teens Who Mastered Fake News: Inside the Macedonian Fake News Complex' *Wired*, 15 February 2017.
8 'Clickbait'—Internet content whose main purpose is to attract attention and encourage visitors to click on a link to a particular web page (Oxford Dictionaries Online).
9 'Post-truth'—'Relating to or denoting circumstances in which objective facts are less influential in shaping public opinion than appeals to emotion and personal belief' (Oxford Dictionaries Online). For a more in-depth discussion of 'post-truth' see Jente Althius and Leonie Haiden (eds), *Fake News: A Roadmap* (Riga and London: NATO Strategic Communications Centre of Excellence and King's Centre for Strategic Communications, 2018), p. 17.
10 'Alt-right'—'a set of far-right ideologies, groups and individuals whose core belief is that 'white identity' is under attack by multicultural forces using 'political correctness' and 'social justice' to undermine white people and 'their' civilization' (Southern Poverty Law Center).
11 M. Allen, 'Kremlin's "Social Media Takeover": Cold War Tactics Fuel Ukraine Crisis' *Democracy Digest* (blog), 10 March 2014; Celestine Bohlen, 'Cold War Media Tactics Fuel Ukraine Crisis', *New York Times*, 10 March 2014, sec. Europe.
12 'Troll'—'a person who starts quarrels or upsets people on the Internet to distract and sow discord by posting inflammatory and digressive, extraneous, or off-topic messages in an online community (such as a newsgroup, forum, chat room, or blog) with the intent of provoking readers into displaying emotional responses and normalizing tangential discussion, whether for the troll's amusement or a specific gain' (Wikipedia).

Discussions on blogs easily spill over into mainstream media, giving information originating from dubious sources the appearance of legitimacy. In August 2016, while a vigorous national debate was underway on whether Sweden should enter a military partnership with NATO, officials in Stockholm suddenly encountered an unsettling problem—a flood of distorted and false information appeared on social media, confusing public perceptions of the issue.[13] The claims were alarming—if Sweden, not a member of NATO, signed the deal, the alliance would stockpile secret nuclear weapons on Swedish soil; NATO could attack Russia from Sweden without government approval; NATO soldiers, immune from prosecution, could rape Swedish women without fear of criminal charges. These stories were all false, but the disinformation was repeated by the traditional news media. As Defence Minister Peter Hultqvist travelled the country to promote the pact in speeches and town hall meetings, he was repeatedly grilled about the bogus stories.

This article begins by examining blogs riddled with mis- and disinformation to identify common characteristics that have been developed as heuristics to detect 'fake-news' blogs. 'Disinformation' is defined as 'the manipulation of information that purposefully aims to mislead or deceive', while 'misinformation' is defined as 'inaccurate information that is the result of an honest mistake or of negligence'.[14] It then goes on to explore the role blogs play in distributing mis- and disinformation. Major social media platforms like Twitter and Facebook are not conducive for agenda setting or framing narrative due to character limitations and the nature of the platform. Blogs do not have any restriction on the number of characters. Bloggers use various modalities for an effective framing, such as images, videos, and audio files. Moreover, there is no risk of getting suspended if the narratives veer into hate speech. Bloggers have the liberty to set the agenda, however blogs lack the social network structure (no friend/follower relations) required to disseminate blog posts via links. Studies have shown that social media platforms are used strategically to coordinate cyber propaganda campaigns.[15] Nine out of ten bloggers have Facebook accounts. Seventy-eight per cent of bloggers use Twitter to promote their content. This percentage is higher for professional and full-time bloggers, almost ninety per cent. In addition to promoting their own content, some bloggers also exploit computer

13 Neil MacFarquhar, 'A Powerful Russian Weapon: The Spread of False Stories', *New York Times*, 28 August 2016, sec. Europe.
14 Althius and Haiden (eds), *Fake News: A Roadmap*.
15 Samer Al-khateeb, Muhammad Hussain, and Nitin Agarwal, 'Leveraging Social Network Analysis & Cyber Forensics Approaches to Study Cyber Propaganda Campaigns', in Tansel Özyer, Sambit Bakshi, and Reda Alhajj, (eds), *Social Network and Surveillance for Society* (Berlin: Springer, forthcoming in 2019).

programmes known as bots, social bots, or botnets, that can massively amplify the dissemination of content via Twitter; this phenomenon has been widely studied.[16] YouTube, another of the fastest-growing social media platforms, is increasingly exploited for behavioural steering, with various production styles catering to specific demographics (such as teens and youth, or even more specific targets such as Somali-speaking men under the age of 35) subjecting the viewers to conspiracy theories, disinformation campaigns, and radicalising ideologies. Prolific linking of YouTube videos in tweets, blogs, or Telegram posts has led to unprecedented convenience in framing narratives, disseminating them widely, and driving online traffic to develop rich discourse.

In addition to content promotion, active media-integration strategies help to artificially boost search rankings. This technique, known as 'link farming', is a well-known search engine optimisation (SEO) strategy. Gaming search engines by prolific linking of blogs across the social-media ecosystem is now part of cyber-influence operations. This paper aims to peel away the layers of the complex media integration strategy to examine the role of cross-media and mix-media strategies (defined below) in conducting disinformation campaigns. By further examining these information flows, we have identified several blog and social media networks responsible for disseminating disinformation in the Baltic States. The paper presents an in-depth examination of such networks, using a social network analysis-based methodology referred to as 'Focal Structure Analysis',[17] to identify the prominent information brokers and leading coordinators of the disinformation campaigns.

The rest of this paper is organised as follows: the Related Work section describes pertinent literature—we review how researchers study disinformation and what techniques they use, and survey the fields that engage in this sort of study, e.g. cyber forensics and social networks; the Data Collection section describes how data is fetched from different sources; the Data Statistics section provides more detail on the attributes of the datasets collected; the Research Methodology section outlines the research questions of the study, explains our approach to answer these questions, and provides a discussion of the findings; and, finally, we summarise the paper with Conclusions.

16 Nitin Agarwal, Samer Al-khateeb, Rick Galeano, and Rebecca Goolsby, 'Examining The Use Of Botnets And Their Evolution In Propaganda Dissemination', *Defence Strategic Communications*, Volume 2 (2017): 87–112.
17 Fatih Şen, Rolf Wigand, Nitin Agarwal, Serpil Tokdemir, and Rafal Kasprzyk, 'Focal Structures Analysis: Identifying Influential Sets of Individuals in a Social Network', *Social Network Analysis and Mining* Volume 6, № 1 (8 April 2016): 17. PAYWALL

Related Work and Terms

Numerous studies have been done on how misinformation and disinformation are spread, sometimes by irresponsible actors and sometimes by malicious ones. Researchers, including Agarwal et al., have studied blogs for a variety of purposes such as identifying the influence of certain bloggers, and how they engage with the audience.[18] Alcott and Gentzkow have studied how blogs are used to disseminate disinformation by linking to platforms such as Twitter, Facebook, and Reddit.[19] Disinformation sharing is amplified through the use of bots. The use of bots has nefarious consequences when consumers of the information they share fail to verify the accuracy of the information and the reliability of the sources. An article published in the *MIT Technology Review* shows that accounts actively spreading false and misleading information are significantly more likely to be bots.[20] Ferrara's study on the use of bots to disseminate information found that they were very effective.[21]

Knowing what is factually correct is not a trivial task, given the barrage of biased, satirical, or conspiracy-theory-riddled stories on social media. The endless deluge of information, sometimes totally irrelevant or junk information, makes it more difficult to stem the flow of mis- and disinformation. Echo chambers quickly emerge, shielding users from divergent opinions, making the problem even worse.[22] There are, however, cases where disinformation has been debunked by such well-known fact-checking websites as snopes.com, politifact.com, and factcheck.org. For instance, a post titled 'No More Child Support After 2017?' was identified as false by snopes.com,[23] and the 'PizzaGate' conspiracy theory was also eventually debunked.[24] However, such community-based fact checking efforts are largely manual, hence severely limited and also subject to personal biases. Automated systems are therefore warranted that can assist in debunking false claims. Few studies have been conducted recently in this domain. Kumar et al. devised methods of classification to determine whether or not a given article

18 Nitin Agarwal, Huan Liu, Lei Tang, and Philip S. Yu, 'Identifying the Influential Bloggers in a Community' in the Proceedings of the 2008 International Conference on Web Search and Data Mining, 207–218.
19 Hunt Allcott and Matthew Gentzkow, 'Social Media and Fake News in the 2016 Election', Working Paper № 23089 (Cambridge, MA: National Bureau of Economic Research, 2017).
20 Emerging Technology from the arXiv, 'First Evidence that Social Bots Play a Major Role in Spreading Fake News', *MIT Technology Review*, 7 August 2017. [accessed 4 June 2018].
21 Emilio Ferrara, 'Measuring Social Spam and the Effect of Bots on Information Diffusion in Social Media' ArXiv 2017.
22 Jonathan Bright, 'Explaining the Emergence of Echo Chambers on Social Media: The Role of Ideology and Extremism', Oxford Internet Institute, University of Oxford, 14 March 2017.
23 Snopes.com, 'FACT CHECK: No More Child Support After 2017?', Kim LaCapria (fact-checker), 20 September 2017, updated 20 April 2018.
24 Snopes.com, 'FALSE: Comet Ping Pong Pizzeria Home to Child Abuse Ring Led by Hillary Clinton', Kim LaCapria (fact-checker), 21 November 2016, updated 4 December 2016.

on Wikipedia is a hoax.[25] Shu et al. developed a data mining model to detect so-called fake news in social media. [26]

Since blogs do not have a social network structure, they use other social media to disseminate their content. In order to discover these often-invisible connections, we use cyber forensic analysis to dive deeper into the problem of disinformation. Cyber forensics is 'the process of acquisition, authentication, analysis, and documentation of evidence extracted from and/or contained in a computer system, computer network, and digital media'.[27] Cyber forensics is a method of collecting metadata used to gain insight into the flow of digital information. Our objective is to understand the patterns associated with dissemination of information posted on blogs, how bloggers build engagement with their audiences, and how they communicate their agendas. We do this by identifying the media approaches used.

The term 'mix-media information dissemination campaign' refers to the use of multiple social media channels to disseminate a narrative. More precisely, an information campaign can be observed on multiple social media sites using similar text, images, audio, and video content. The content may not be identical on the various social media channels used, but pertains to one particular information campaign.

A 'cross-media information dissemination campaign' orchestrates the use of specific media channels. More precisely, the information is hosted on a website (e.g. a blog site or video on a YouTube channel) and is widely distributed through other social media channels that provide established social network structures, such as Twitter and Facebook.

Until recently, cyber forensic tools, or 'the practice of identifying, preserving, extracting, analysing, and presenting legally sound evidence from digital media such as computer hard drives',[28] have been primarily used by government agencies, but as the need for tracking such data has grown, many cyber forensic tools, such as Maltego, have become available to the public.[29] Cyber forensics studies can also

25 Srijan Kumar, Robert West, and Jure Leskovec 'Disinformation on the Web: Impact, Characteristics, and Detection of Wikipedia Hoaxes' in *Proceedings of the 25th International Conference on World Wide Web*, pp. 591–602, International World Wide Web Conferences Steering Committee, 2016.
26 Kai Shu, Amy Sliva, Suhang Wang, Jiliang Tang, and Huan Liu. 'Fake News Detection on Social Media: A Data Mining Perspective', *SIGKDD Explorations* 19, Issue 1(June 2017): 22–36.
27 Digambar Povar and V. K. Bhadran, 'Forensic Data Carving' in *International Conference on Digital Forensics and Cyber Crime, Second International ICST Conference*, Abu Dhabi, United Arab Emirates, 4–6 October 2010, Revised Selected Papers, ed. Ibrahim Baggili (Berlin: Springer, 2011): 137–48. PAYWALL.
28 Snopes, 'FALSE: Comet Ping Pong Pizzeria'.
29 Paterva.com, Company Page. [Accessed 4 June 2018]

take a variety of perspectives. Sabillion et al. have subjected multiple environments in digital ecosystems to digital forensic investigations from an information security and privacy standpoint,[30] while Harichandran et al. conducted a survey-based analysis from a computer security background.[31] Our study is unique in that we have conducted our analyses using open source information from social media to identify the key actors disseminating mis- and disinformation.

Data Collection

Data collection is a critical task for analysing information flow in social media. As blogs are becoming virtual town halls[32] that shape discourse, and have already become the preferred mode of carriage for all other digital media, we explore blogs in greater depth to understand the role they play in distributing mis- and disinformation.

Blog sites differ from one another in structure, so collecting data about them necessarily involves sophisticated, tailor-made techniques. 'Crawling' blog data requires high attention to detail and cannot be fully automated. However, other social media platforms usually have a data extraction limit, a rate limit, or both, so there is an advantage to extracting data from blogs.

These are some of the challenges we faced during the data crawling process:

- **No application programming interface (API) to collect data**

  API in its simplest form is an interface that allows an application to talk to another application through simple commands. There is no blog data API that allows the user to make a request (i.e. crawl blog posts) and get a required response (i.e. blog posts data). There were blog indexing services such as BlogPulse, Blogdex, and Technorati that provided limited blog data APIs. But these efforts have been discontinued, making blog data collection a challenging task.

30 Regner Sabillon, Jordi Serra-Ruiz, Victor Cavaller, and Jeimy J. Cano, 'Digital Forensic Analysis of Cybercrimes: Best Practices and Methodologies', *International Journal of Information Security and Privacy*, Volume 11, № 2 (2017): 25–37. PAYWALL

31 Vikram S. Harichandran, Frank Breitinger, Ibrahim Baggili, and Andrew Marrington, 'A Cyber Forensics Needs Analysis Survey: Revisiting the Domain's Needs a Decade Later', *Computers & Security* 57 (2016): 1–13.

32 '**Town hall meetings**, also referred to as **town halls** or **town hall forums**, are a way for local and national politicians to meet with their constituents, either to hear from them on topics of interest or to discuss specific upcoming legislation or regulation. During periods of active political debate, town halls can be a locus for protest and more active debate.' (Wikipedia).

DOI 10.30966/2018.RIGA.4.6.

- **Changing blog structure**

  Dealing with blogs is like shooting at a moving target. Blog site owners can change the blog site structure at any time, and a crawler trained for one structure does not work for another. So the effort of training the crawler must be repeated for the new structure. Additionally, each blog requires its own parser to crawl the data. As in linguistics, in programming to parse is to separate something into its component parts, so that these parts can be stored and manipulated. As the blog sites have different structures (syntax), different parsers are required for different blogs as a parser trained for one blog may not work at all for others.

- **Noise**

  Regardless of how well a crawler is trained, 'noise' is always crawled. Social media plugins (Facebook share plugins, Twitter share plugins) and advertisements from the blog site could be crawled as part of noise. Even though the crawler is well trained, some JavaScript elements are often crawled as part of the blog data. This is considered noise as it does not constitute any useful information. Noise is later removed as part of the data cleaning process.

- **No standardisation**

  While we collect blog data, we parse important attributes for analysis. One such attribute is date. While we work with date field extraction from blog posts, we notice that these differ in format from blog site to blog site, implying that there is no single standard followed here. This adds to the workload of converting discrepancies of this kind to a standard format for further study.

- **No automation**

  The blog crawling process is not fully automated. Even the most intelligent and careful parsing may capture some noise. A human researcher must be in the loop to identify and eliminate noise.

- **Limitations of the Web Content Extractor (WCE)**

  We use the WCE tool for blog data collection. With the help of this software <http://www.newprosoft.com/> we train the crawler to extract data from different blog sites efficiently. To train the crawler, we first provide the starting

or 'seed URL'. Then, we train the crawler to navigate to all blog posts linked to the seed URL. While doing this, we sometimes encounter noise and observe that the WCE fails in crawling dynamic pages loaded with JavaScript.[33]

We examined several blogs and identified a few common attributes among them such as title, date of posting, author/blogger, actual post, number of comments, and permalink. Later, we collected and indexed all blog content from the various sources to our Blogtrackers database. This tool can be accessed at <http://blogtrackers.host.ualr.edu/>. We surveyed over 300 blogs, including: a 'fake news' or disinformation dataset from Kaggle.com, disinformation blogs identified by Dr Melissa Zimdars, a professor at Merrimack College,[34] blogs containing disinformation regarding the Baltic States, and blogs containing disinformation regarding NATO activities and exercises. To crawl these different sources, we set up crawler(s) for each blog site to extract all required attributes. The main steps in crawling data from a blog site are: (1) exploring the blog site; (2) crawling the blog site; and (3) cleaning and storing the data in a database for retrieval and analysis. Figure 1 below represents the data crawling process for blogs.[35]

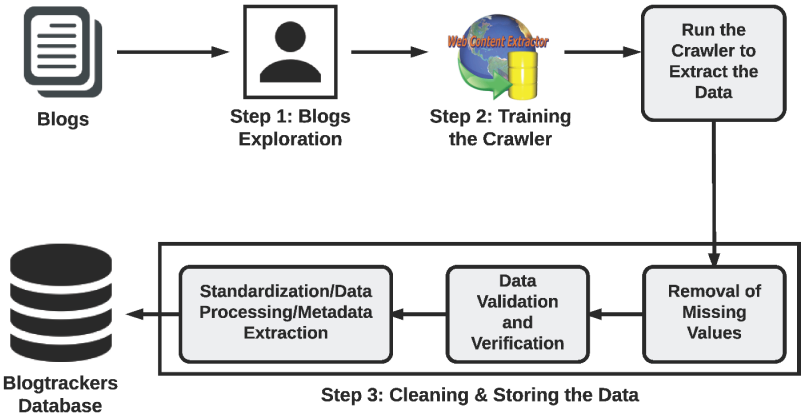

Figure 1. Data collection process for blogs

33 Muhammad Nihal Hussain, Adewale Obadimu, Kiran Kumar Bandeli, Nooman Mohammad, Samer Al-khateeb, and Nitin Agarwal, 'A Framework for Blog Data Collection: Challenges and Opportunities', The IARIA International Symposium on Designing, Validating, and Using Datasets (DATASETS 2017), June 2017.
34 Merrimack College, Dr Melissa Zimdars
35 Muhammad Nihal Hussain, Saaduddin Ghouri Mohammad, and Nitin Agarwal, 'Blog Data Analytics Using Blogtrackers', International Conference on Social Computing, Behavioral-Cultural Modeling & Prediction and Behavior Representation in Modeling and Simulation, Washington DC, July 2017.

Next we added 'data augmentation'. This additional step facilitates the use of cyber forensics to collect metadata such as language, location, web traffic tracker codes, and other data. We used the cyber forensics tool Maltego to extract the metadata. During cyber propaganda campaigns, the extracted metadata made it possible to discover the information we used to track IP addresses and to identify the groups of users or entities that were working together to disseminate disinformation, as well as to discover hidden connections among the blogs.

Data Statistics

Primarily, the data collected is from four diverse sources. The descriptions and statistics associated with these four types of datasets are as follows:

A. A Disinformation dataset from Kaggle.com

This dataset was obtained from the Kaggle website using the link <https://www.kaggle.com/mrisdal/fake-news>. Data was collected for a period of 30 days from 2016.10.07 to 2016.11.07. It has 20 attributes, 2,236 bloggers, and 12,999 posts from 244 blogs. Some of the key attributes in this dataset are—domain name, site_url, author, post title, text, published date, language, comments, replies_count, shares, and likes.

B. Dr Zimdars' list of Disinformation blogs

These blog sites are available at <http://bit.ly/2ezvFbV>. This dataset had 26 blogs, 971 bloggers, 116,667 posts, and 79 attributes at the time of this study. The list has now grown to over 1000 blogs thanks to an international community-wide effort. This dataset was collected from 2004.01.02 to 2017.05.04, and its key attributes are—blog site name, blogger, blog post title, blog post, posting date, location, and language. Tables 1 and 2 show the statistics for this dataset, describing the location and language distribution of the blogs.

Location	Blog Posts	Blogs
United States	79768	21
Not Available	20001	4
Great Britain	16898	1

Table 1. Location distribution for the 26 disinformation blogs

DOI 10.30966/2018.RIGA.4.6.

Language	Blog Posts	Blogs
English	115689	26
Not Available	926	5
Danish	9	2
Unknown	9	2
French	7	4
German	6	4
Guerrero	3	1
Italian	3	2
Finnish	2	2
Spanish	2	1
Afrikaans	1	1
Catalan	1	1
Dutch	1	1
Estonian	1	1
Hiligaynon	1	1
Latin	1	1
Maltese	1	1
Norwegian	1	1
Romanian	1	1
Russian	1	1
Swedish	1	1

Table 2. Language distribution for the 26 disinformation blogs

C. Blogs containing disinformation regarding the Baltic States

Subject matter experts identified this set of blogs suspected of disseminating disinformation, many of them originating from Latvia, Estonia, or Lithuania. We crawled 16,667 blog posts from 21 blogs with 728 bloggers. This dataset was collected from 2005.03.24 to 2017-02-09, and the key attributes are—blog site name, blogger, blog post title, blog post, posting date, location, and language. Tables 3 and 4 provide the statistics for this dataset, describing the location and language distribution of the blogs respectively.

Location	Blog Posts	Blogs
Lithuania	6595	2
United States	3156	3
Estonia	2592	7
Germany	2156	5
Latvia	1976	3
Netherlands	192	1

Table 3. Location distribution for the 21 Baltic States' blogs

Language	Blog Posts	Blogs
Lithuanian	6590	2
Russian	4599	7
Latvian	3793	10
Estonian	809	6
English	738	8
Unknown	122	5
Not Available	15	1
Bulgarian	1	1

Table 4. Language distribution for the 21 blogs about the Baltic States

D. Blogs containing disinformation regarding NATO exercises and other activities

This dataset contains blogs containing disinformation that posted news during various exercises conducted by NATO, such as the Trident Juncture 2015, Brilliant Jump 2016, and Anakonda 2016. The dataset was collected using the Blogtrackers tool between 1993.02.28 and 2017.08.06. It has 70 blogs, 3641 bloggers, 118,908 posts, and 79 attributes. The key attributes are—blog site name, blogger, blog post title, blog post, posting date, location, and language. Tables 5 and 6 show the statistics for this dataset, describing the location and language distribution of the blogs.

Location	Blog Posts	Blogs
Germany	1788	4
France	228	2
Ireland	26	1
Netherlands	7572	3
Poland	16	1
Serbia	36	1
Russia	5419	4
Slovakia	285	1
Ukraine	391	1
United States	103146	51
Zimbabwe	1	1

Table 5. Location distribution for the 70 NATO exercise blogs

Language	Blog Posts	Blogs
English	101776	64
Polish	6107	10
German	2598	15
Unknown	1893	30
Croatian	1851	1
Spanish	1565	15
Romanian	641	5
Russian	623	11
Italian	590	9
Not Available	369	11
French	316	16
Greek	181	2
Catalan	115	5
Arabic	102	4
Ukrainian	46	3
Albanian	43	1

Czech	12	4
Danish	10	7
Afrikaans	9	2
Slovak	8	2
Turkish	8	3
Basque	6	1
Finnish	6	3
Korean	5	1
Portuguese	4	3
Dutch	3	2
Hungarian	3	2
Serbian	3	2
Swahili	3	2
Welsh	3	2
Central Mam	2	1
Faroese	2	1
Indonesian	2	2
Latvian	1	1
Maltese	1	1
Tagalog	1	1

Table 6. Language distribution for the 70 NATO exercise blogs

The attributes from the four datasets are summarised in Table 7.

Dataset	Number of Blogs	Bloggers	No of Posts	Attributes
Disinformation dataset from Kaggle.com	244	2236	12999	20
Disinformation blogs identified by Prof. Melissa Zimdars	26	971	116667	79
Blogs containing disinformation regarding Baltic States	21	728	16667	79
Blogs containing disinformation regarding NATO exercises and activities	70	3641	118908	79

Table 7. Details of the four datasets summarised

Research Methodology

This barrage of information on social media illustrates the challenges encountered in studying disinformation campaigns, challenges caused by the high volume and variety of the data. To classify such information as true or false poses big questions for researchers in the field. By the time a blog post hosting disinformation is detected and appropriate action is undertaken to stop it, the damage has already been done. Therefore, we study dissemination to understand the strategies and tactics used by those who operate such blogs. To this end, we formulated the following research questions:

- Research Question 1: What are the typical characteristics of mis- and disinformation-riddled blogs?
- Research Question 2: Can we track the origins of the content, such as memes, images, etc. appearing in these blogs?
- Research Question 3: What strategies (mix-media or cross-media) and media sites are commonly used for disseminating such content?
- Research Question 4: Is there a way to track how far a narrative antagonistic to NATO travels, and how effective the content has been in terms of comments, likes, and other markers, relative to other narratives?

The aforementioned research questions are answered during different phases in the proposed methodology, as shown in Figure 2 below. To answer RQ1, we collected data from different disinformation-riddled blogs, studied their characteristics, and proposed guidelines for identifying disinformation. To answer RQ2, we studied the provenance of the information on these blogs, for example, by tracking the sources using reverse image search, memes, web links, and hashtags. To answer RQ3, we studied the strategies commonly used to disseminate disinformation—the mix-media and cross-media approaches. Finally, to answer RQ4, we extended the work done on RQ1. The blog network was identified through the characteristics extracted from the blogs; we then used the Focal Structure Analysis (FSA) tool, an aspect of social network analysis, to identify the key disinformation actors or groups.

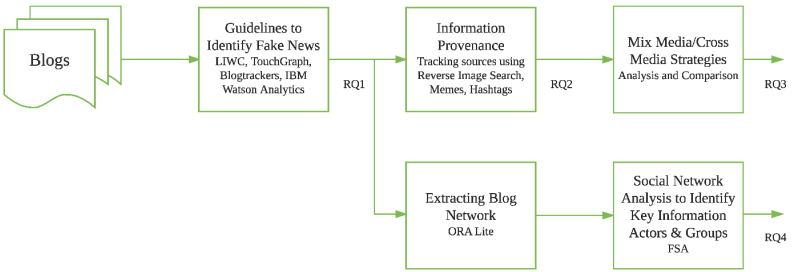

Figure 2. Proposed research methodology

RQ1: Typical Characteristics of Disinformation-riddled Blogs

We started the experiment by loading the datasets into IBM Watson Analytics to identify the data attributes, or the structure of the data, so as to detect key patterns that reveal the typical characteristics of disinformation-riddled blogs. To this end, we considered websites that frequently shared posts—their titles, likes, replies, shares, and comments. Then, we used tools such as TouchGraph SEO Browser,[36] LIWC (Linguistic Inquiry and Word Count),[37] and Alchemy API[38] to dive deeply into the problem space. Later, we used Blogtrackers to analyse trends in keywords of interest, to identify unusual patterns in bloggers' posting frequency and changes in topics of interest to the blogging community

36 TouchGraph, SEO Keyword Graph Visualisation. http://www.touchgraph.com/seo [Accessed 4 June 2018]
37 LIWC | Linguistic Inquiry and Word Count, Landing Page. http://liwc.wpengine.com [Accessed 4 June 2018]
38 <https://www.ibm.com/watson/alchemy-api.html>

in real time, to track trends in positive and negative sentiment, to estimate a blogger's influence in the blogosphere, and to understand the outlinks used to disseminate content through social media channels.

Based on our own observations and on those compiled by other experts, we provide a set of heuristics to identify blogs potentially riddled with disinformation:

1. **Pay attention to the *contact us* section of a page to validate and verify site authors.**

   Misinformation-riddled blogs do not provide any real information about the website author. See: <http://abcnews.com.co/>

2. **Do not just read the *headline*; instead examine the *body content* to discover more details of the story.**

   A headline, such as 'Obama Signs Executive Order Declaring Investigation into Election Results; Revote Planned for Dec. 19th—ABC News' is provocative, but the content of the story may reveal information useful to determining a post's veracity.

3. **Pay close attention to URLs, sources, *images,* and to the editorial standards of the writing.**

   The URL <www.bloomberg.ma> is used to imitate the well-known site <www.bloomberg.com>.

4. **Always crosscheck the story with *fact-checking websites*, such as snopes.com, factcheck.org, or politifact.com to verify the credibility of the story.**

   Is the blog post 'The Amish In America Commit Their Vote To Donald Trump; Mathematically Guaranteeing Him A Presidential Victory—ABC News' credible? Snopes says no.

5. **Search for the post in *well-known search engines* such as Google, Bing, or Yahoo. If the same post or content is repeated on many sites, this indicates the use of the mix-media approach to narrative dissemination.**

   The disinformation blog post 'Obama Signs Executive Order Declaring Investigation Into Election Results; Revote Planned For Dec. 19th—ABC News' was shared on many sites across the web.

6. **Check if the article has been *previously published* and if it is being reused to affect perceptions about an event or specific actions.**

The blog post 'Muslims BUSTED: They Stole Millions In Govt Benefits' published in 2016 reused an image from 2013.

7. **Check to see if the post is *disturbing or controversial;* fake stories are often embedded in such posts.**

   The blog post 'EU NATO Commit Adultery, Prince Charles Saudi Trade & More' presents disturbing information. Disinformation narratives are framed in just such posts.

8. **Check to see if the post has any likes, replies, or comments. This will indicate how interested the readers are in the given story, e.g. agree or disagree with it. The *sentiment* can be used to infer this.**

   The blog post 'NASA Confirms—Super Human Abilities Gained' received numerous comments, many of them debunking this false story. Sentiment analysis on commentary provides a clear understanding of a post's reception.

We conducted a survey to evaluate the efficacy of these eight criteria for predicting the veracity of a blog. We randomly selected 96 blog sites featuring mis- or disinformation and asked survey participants to rate the effectiveness of each criterion.

After collecting the survey data, we constructed a stacked bar where the Y-axis denotes the eight criteria, and the X-axis (representing values 0%–100%) indicates the confidence of the participants of the veracity of the 96 blog sites, classified as low, medium, or high. Figure 3 clearly shows that the best criterion is a blog site's use of mix-media and/or cross-media strategies for disseminating content. The next best criterion is use of a fact-checking website.

Next, we present some empirical observations vis-à-vis disinformation heuristics on the disinformation dataset collected from Kaggle.com. Incidentally, most of the blog posts had no or very few comments; this might imply that these stories are disseminated, but not much discussed. We also discovered that, during the US elections, many posts were designed to direct the greatest number of readers to non-factual stories with the intention of biasing the electorate. For example, 12,468 of 12,999 (96%) of posts had zero 'likes' and 12,304 of 12,999 (94%) posts had zero 'replies'. We also observed that the majority of the stories originated from a particular set of domains usually reported as containing false info.

DOI 10.30966/2018.RIGA.4.6.

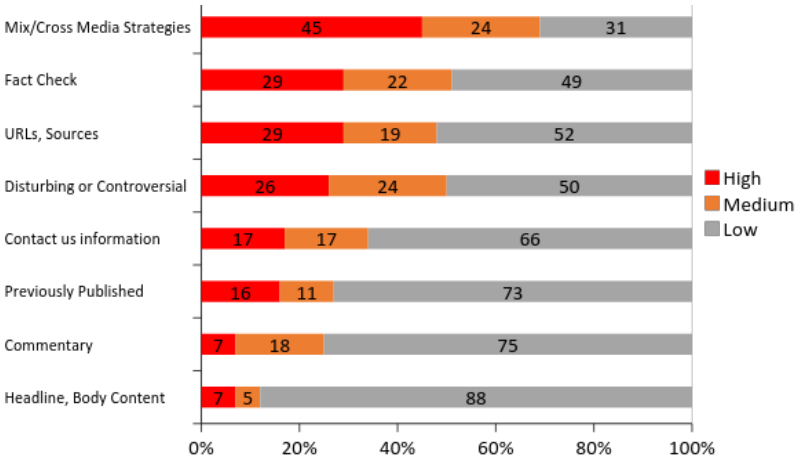

Figure 3. The effectiveness of each of the eight criteria in identifying disinformation-riddled blog

We found that in many cases that the 'contact us' page did not provide any real contact information, or it redirected readers to another website, usually a social media site, as shown in Figure 4 below: ABC NEWS at <http://abcnews.com.co/>. The <http://cnn.com.de/> website closely mimics the CNN News website <http://www.cnn.com/> using the CNN logo, its website structure, etc. However, the impostor site is riddled with false stories and conspiracy theories, but when articles are re-posted on Facebook, they bear the CNN.com logo and appear to be published by the authentic site. This deception strategy is highly effective and is further discussed in the next two sections. On a final note, the indicators used for answering RQ1 can be fine-tuned by using advanced language analytics. This remains a task for the future.

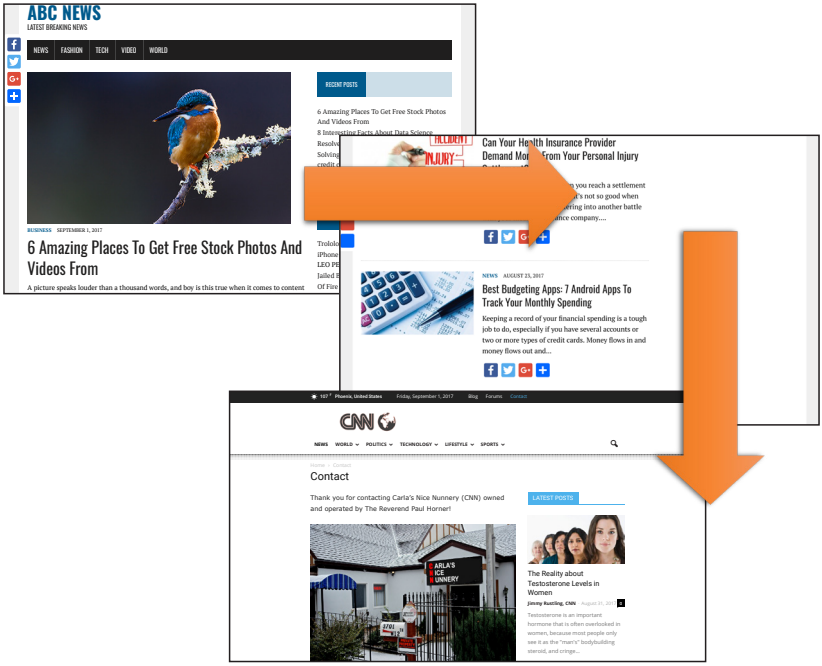

Figure 4. Contact link redirects to another website

RQ2: Tracking the Origins of Blog Content

In this section, we track the origins of memes, images, and other features appearing on selected blogs sites. We begin by searching for a given image URL using 'reverse image search' on Google Images to identify other sources using the same image. We found that many of the images used were not unique for each article, and often not relevant to the context of the article. The same image was reused with multiple different narratives, as shown in Figure 5. Images lend credibility to a narrative and are effective in fabricating perception. Using images or videos to frame narratives is highly effective, exploiting multiple modalities to influence thinking.

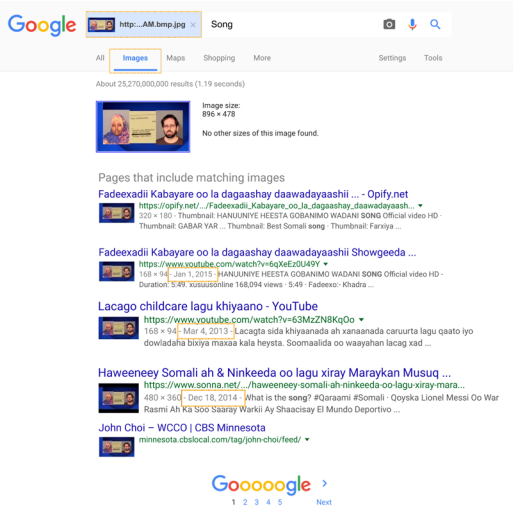

Figure 5. Reverse image search shows the same image used to illustrate multiple narratives

We observed a pattern frequently used to disseminate malicious content: purposefully generated content is posted to a blog; hashtags with embedded links serve to connect the post to high-traffic social media channels. This is known as shared or referenced using. Figures 6(a) and 6(b) below depict this pattern.

Figure 6(a). Post using hashtags with embedded hyperlinks to direct traffic to Twitter

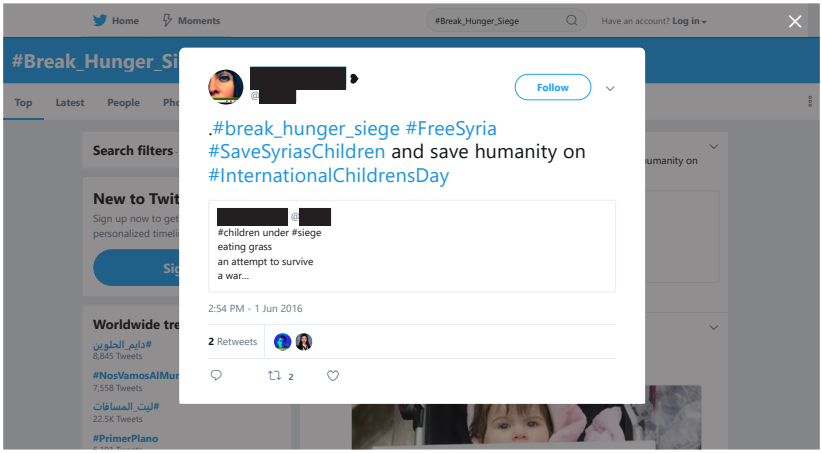

Figure 6(b). Twitter post using hyperlink-hashtags

RQ3: Mix-media vs Cross-media Approaches

To recap, when the designers of an information dissemination campaign post similar or identical texts, images, audio, and video content to multiple social media platforms we call this a mix-media campaign; when content is hosted on a single platform, such as a blog or a YouTube channel, but is distributed widely across multiple social media channels, we call this a cross-media campaign. The mix-media strategy is more work-intensive and tailor-made, while the cross-media relies more on strategic link placement. We hoped to discover which of these strategies is most commonly used to disseminate bogus content and if we can use our knowledge of these strategies to identify the other media sites in a dissemination chain?

We first investigated the use of the mix-media approach. The example in Figure 7 shows that the story 'Towards a Renewed Imperialist Intervention in Libya? Anti-NATO Forces Retake Areas in Southern Libya' was disseminated on multiple sites—facebook.com, oroom.org, twitter.com, globalresearch.ca, hotnews.ro, and workers.org.[39]

39 This story was marked as a conspiracy theory by the application bsdetector.tech

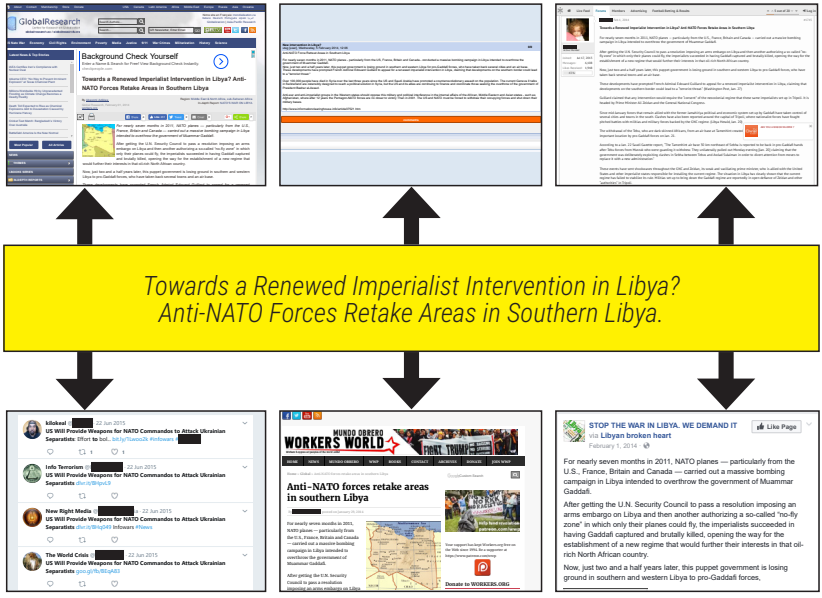

Figure 7. The mix-media strategy: disseminating disinformation on different websites.

Next, we examined the cross-media information dissemination approach. This strategy was observed to good effect in our dataset. Many sites that shared links to specific social media channels such as Twitter, Facebook, and Reddit. For instance, a blog site named 'globalresearch.ca' published a post entitled 'US Will Provide Weapons For NATO Commandos to Attack Ukrainian Separatists' with a link: http://bit.ly/2ewVTg7. This post was shared on a variety of social media channels—Twitter: http://bit.ly/2xEQxnU, Pinterest: http://bit.ly/2x02sQ0, and Facebook: http://bit.ly/2wrIhZD, as depicted in Figure 8. This clearly indicates the use of the cross-media approach.

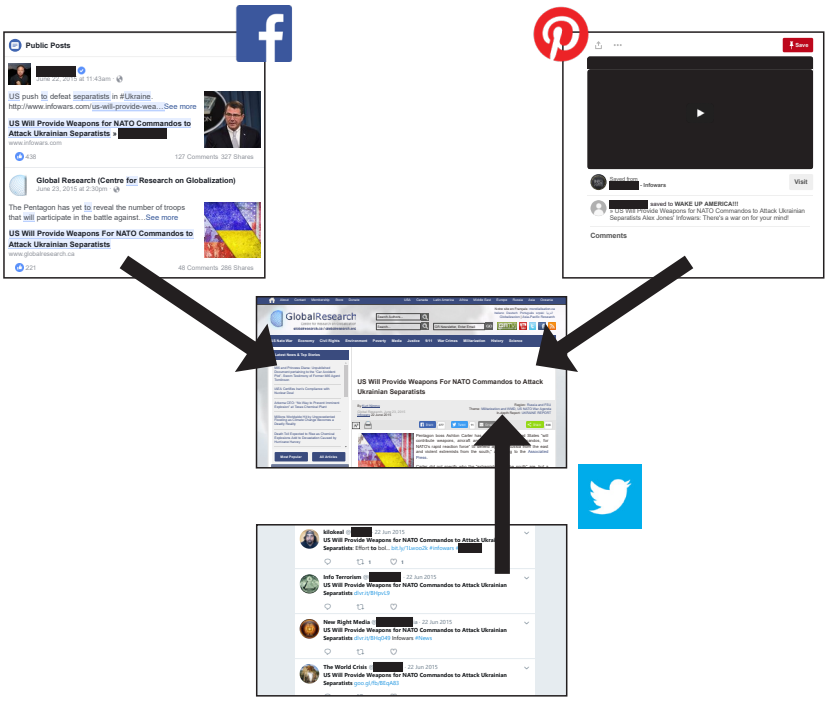

Figure 8. The cross-media strategy for disseminating disinformation on different social media channels

RQ4: Tracking How Far an Antagonistic Narrative Travels

We wanted to discover if we could track how far antagonistic narratives travel and how effective such narratives are, what kinds of reactions they received. We also analysed the blog networks to identify the blog clusters fundamental to disseminating these narratives. We considered reactions such as likes, comments, retweets, and shares to discover how widely a post had been circulated, and to uncover the media integration strategies used.

Readers may not realise that social plugins from Twitter, Facebook, and Reddit can be embedded in the 'like' feature of blogs. When a blog post is 'liked', this reaction shows up simultaneously on all linked platforms. For instance, on 18

September 2016 the blog site 21stcenturywire.com published the post 'Syria: No "Dusty Boy" Outrage for 7 yr old Haider, Sniped by NATO Terrorists in Idlib Village of Foua'. This post received 65 comments and was shared on other social media. Its Twitter repost received 19 retweets, 5 likes, and 2 replies, while on Facebook it received 6 reactions, 3 comments, and 2 shares. This disinformation effort achieved modest success. 21stcenturywire.com published another post on 27 September 2016 entitled 'EU NATO Commit Adultery, Prince Charles Saudi Trade & More', which again was factually incorrect. As before, we tracked how this post travelled; however, it received no comments. The article was shared on Twitter, but it received only 1 retweet, 1 like, and 0 replies. The same post was also shared on Facebook, where it received 27 reactions, 1 comment, and 11 shares. But, all the shares came from 21stcenturywire.com. No other Facebook group posted this article. This one had no traction.

Next, we analysed the effect of blog networks on content dissemination. Unlike other social media platforms, blogs do not have a social network structure—there is no follow-follower relationship among blogs. However, information flow can be tracked based on who links to whom. We examined hyperlinks to extract information regarding the network of the blogs generating disinformation regarding the Baltic States using ORA-Lite (Organisation Risk Analyser)[40] to visualise the network, as depicted in Figure 9 below. The network contains 21 blogs (red nodes) and 2,321 hyperlinks (blue nodes). Further analysis of the blog network helped identify five blogs (out of the set of 21 blogs) that were the most resourceful. In other words, these five blogs had the highest number of hyperlinks, i.e. links to other websites. Furthermore, these five blog sites were the most exclusive in resources, i.e. they linked to websites that no other blogs did. Ten out of the 2,321 hyperlinks were the most shared and the most exclusive, i.e. these hyperlinks were shared profusely but only by a handful of blogs. Most of these hyperlinks have a domain suffix from the Baltic nations, i.e. '.ee' for Estonia, '.lv' for Latvia, and '.lt' for Lithuania. The exclusivity of resources helps in identifying unique information sources and also hints at a coordinated information campaign. This is the subject of our investigation presented in the next section.

40 <http://www.casos.cs.cmu.edu/projects/ora/software.php>

DOI 10.30966/2018.RIGA.4.6.

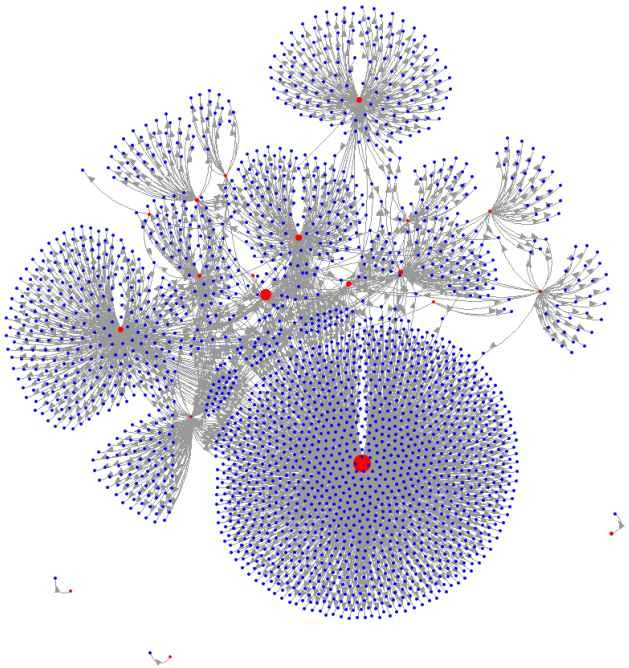

Figure 9. Network of blog sites and shared hyperlinks

The network contains 21 blogs (red nodes) and 2,321 hyperlinks (blue nodes). The size of a node is proportional to the number of shared hyperlinks (i.e. out-degree centrality). Edge thickness is proportional to the number of times a blog site shared a hyperlink.

Focal Structure Analysis

The exclusivity of resource sharing by a few blogs hints at a coordinated information campaign. The resources, in this case—hyperlinks, that are exclusively shared by a set of blogs. To get a clearer picture, we constructed a diagram of the blog network based on shared hyperlinks, depicted in Figure 10. The network is fully connected. Each blog connects to all other blogs, i.e. they form a clique when it comes to information sharing. This suggests that every blog in this set shares the same hyperlinks. This confirms our conjecture

that there is intensive campaign coordination among these blogs. Further investigation is required to discover if these blogs belong to or are controlled by the same individual or by a group.

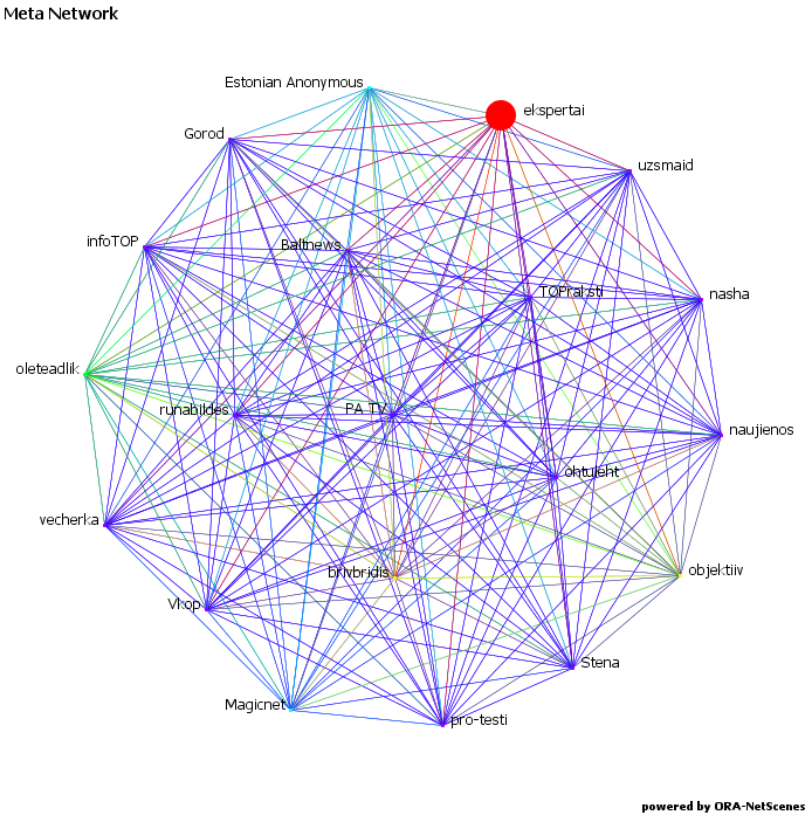

Figure 10. A network of blogs based on commonly shared hyperlinks

To answer the second part of the question—to see how effective the content has been, in terms of comments and likes relative to other narratives—we present an example where commentary forms a persuasive dimension. The research is grounded in exemplification theory and demonstrates how exemplified accounts

Figure 11. An exemplified account of comments shaping discourse to form a persuasive dimension

in user comments on a news story may influence audience perception.[41] Exemplification theory refers to the idea that a single event, or story about that event, is taken to be representative of the set of similar events—it becomes an examplar. Exemplars are neither consumed nor recalled equally, rather those that are engaging to us are often also taken to be representative, even if this is not the case. In other words, media that include engaging, meaningful, attractively illustrated exemplars are more convincing and memorable than well- argued and scrupulously documented news stories.

Guardianlv.com posted the fictitious article 'NASA Confirms—Super Human Abilities Gained'. What could we learn from the 327 content-related comments? The majority of commentators debunked the false information contained in the article, forming a persuasive dimension. Figure 11 illustrates how the 'exemplified accounts' in users' comments can influence perception by developing a

41 Patric R. Spence, Kenneth Lachlan, Timothy Sellnow, Robert G. Rice, and Henry Seeger, 'That Is So Gross and I Have to Post About It: Exemplification Effects and User Comments on a News Story', *Southern Communication Journal*, Volume 82, № 1 (2017): 27–37. PAYWALL

persuasive discourse.[42] The content was linked to other websites or pages (Facebook fan pages) and shared 50 times on other channels, generating further discussion, as illustrated in Figure 11.

Conclusion

Several hundreds of thousands of blog posts, spread across 15 countries, and written in 42 languages, were analysed in this study, which was conducted from 28 February 1993 to 6 August 2017. We observed intriguing patterns of information diffusion used within the social media ecology to spread disinformation. While narratives are framed on blogs, they are disseminated through a variety of other social media channels. The methodology presented here contributes to the understanding of how such disinformation spreads. We documented typical characteristics of disinformation-riddled blogs, using a variety of information analysis tools, and presented guidelines to help identify disinformation on blogs. We demonstrated how to track the origins of the content on blogs with the help of memes, hashtags, links (URLs), and originally-published content. We found that both mix-media and cross-media strategies were commonly used in disseminating the content. We studied how disinformation spreads across different social media platforms using these approaches. We also studied how far an antagonistic narrative travels by analysing blog networks. By applying focal structures analysis, we succeeded in identifying a massively coordinated information campaign among the blogs. And by leveraging exemplification theory, we demonstrated that commentary on blogs can form a persuasive dimension to discourse. Through this study, we have sought to understand the social media ecology and how it supports disinformation dissemination, generating disinformation, weaponising narratives, and conducting propaganda campaigns, so that actions can be taken towards developing countermeasures to stem the tide of fakery.

Acknowledgments

This research is funded in part by the US National Science Foundation (IIS-1636933, ACI-1429160, and IIS-1110868), the US Office of Naval Research (N00014-10-1-0091, N00014-14-1-0489, N00014-15-P-1187, N00014-16-1-2016, N00014-16-1-2412, N00014-17-1-2605, N00014-17-1-2675), the US Air Force Research Lab, the US Army Research Office (W911NF-16-1-0189), the

42 Dolf Zillmann, 'Exemplification Theory of Media Influence' Media Effects: Advances in Theory and Research 2 (2002): 19–41. http://academic.csuohio.edu/kneuendorf/quillin/Gerbneretal%20growing%20up%20with%20television%202002.pdf

US Defense Advanced Research Projects Agency (W31P4Q-17-C-0059), and the Jerry L. Maulden/Entergy Fund at the University of Arkansas at Little Rock. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding organisations. The researchers gratefully acknowledge the support.

Bibliography

Agarwal, Nitin, Samer Al-khateeb, Rick Galeano, and Rebecca Goolsby, 'Examining The Use Of Botnets And Their Evolution In Propaganda Dissemination', *Defence Strategic Communications*, Volume 2 (2017): 87–112.

Agarwal, Nitin, Huan Liu, Lei Tang, and Philip S. Yu, 'Identifying the Influential Bloggers in a Community' in the Proceedings of the 2008 International Conference on Web Search and Data Mining, 207–218.

Al-khateeb, Samer, Muhammad Hussain, and Nitin Agarwal, 'Leveraging Social Network Analysis & Cyber Forensics Approaches to Study Cyber Propaganda Campaigns', in Tansel Özyer, Sambit Bakshi, and Reda Alhajj, (eds), *Social Network and Surveillance for Society* (Berlin: Springer, forthcoming in 2019).

Allcott, Hunt and Matthew Gentzkow, 'Social Media and Fake News in the 2016 Election', Working Paper № 23089 (Cambridge, MA: National Bureau of Economic Research, 2017).

Allen, M., 'Kremlin's "Social Media Takeover": Cold War Tactics Fuel Ukraine Crisis' *Democracy Digest* (blog), 10 March 2014;

Barthel, Michael and Amy Mitchell, 'Americans' Attitudes About the News Media Deeply Divided Along Partisan Lines', *Pew Research Center's Journalism Project* (blog), 10 May 2017.

Blogging statistics: blog posts written today, is a real-time counter showing the global activity on wordpress bloggers. [Accessed 4 June 2018].

Bohlen, Celestine, 'Cold War Media Tactics Fuel Ukraine Crisis', *New York Times*, 10 March 2014, sec. Europe.

Bright, Jonathan, 'Explaining the Emergence of Echo Chambers on Social Media: The Role of Ideology and Extremism', Oxford Internet Institute, University of Oxford, 14 March 2017.

Ferrara, Emilio, 'Measuring Social Spam and the Effect of Bots on Information Diffusion in Social Media' ArXiv 2017.

Harichandran, Vikram S., Frank Breitinger, Ibrahim Baggili, and Andrew

Marrington, 'A Cyber Forensics Needs Analysis Survey: Revisiting the Domain's Needs a Decade Later', *Computers & Security* 57 (2016): 1–13.

Hindman, Matthew, *The Myth of Digital Democracy* (Princeton, NJ: Princeton University Press, 2008).

Hussain, Muhammad Nihal, Saaduddin Ghouri Mohammad, and Nitin Agarwal, 'Blog Data Analytics Using Blogtrackers', International Conference on Social Computing, Behavioral-Cultural Modeling & Prediction and Behavior Representation in Modeling and Simulation, Washington DC, July 2017.

Hussain, Muhammad Nihal, Adewale Obadimu, Kiran Kumar Bandeli, Nooman Mohammad, Samer Al-khateeb, and Nitin Agarwal, 'A Framework for Blog Data Collection: Challenges and Opportunities', The IARIA International Symposium on Designing, Validating, and Using Datasets (DATASETS 2017), June 2017.

Jost, Kenneth and Melissa J. Hipolit, *Blog Explosion: Are Blogs a Passing Fad or a Lasting Revolution?*, *CQ Researcher* 16 (2006), Issue 22 (9 June). PAYWALL.

Kumar, Srijan, Robert West, and Jure Leskovec 'Disinformation on the Web: Impact, Characteristics, and Detection of Wikipedia Hoaxes' in *Proceedings of the 25th International Conference on World Wide Web*, pp. 591–602, International World Wide Web Conferences Steering Committee, 2016.

LIWC | Linguistic Inquiry and Word Count, Landing Page. [Accessed 4 June 2018]

MacFarquhar, Neil, 'A Powerful Russian Weapon: The Spread of False Stories', *New York Times*, 28 August 2016, sec. Europe.

MIT Technology Review, Emerging Technology from the arXiv, 'First Evidence that Social Bots Play a Major Role in Spreading Fake News', 7 August 2017. [accessed 4 June 2018].

Paterva.com, Company Page. [Accessed 4 June 2018]

Povar, Digambar and V. K. Bhadran, 'Forensic Data Carving' in *International Conference on Digital Forensics and Cyber Crime, Second International ICST Conference*, Abu Dhabi, United Arab Emirates, 4–6 October 2010, Revised Selected Papers, ed. Ibrahim Baggili (Berlin: Springer, 2011): 137–48. PAYWALL.

Sabillon, Regner, Jordi Serra-Ruiz, Victor Cavaller, and Jeimy J. Cano, 'Digital Forensic Analysis of Cybercrimes: Best Practices and Methodologies', *International Journal of Information Security and Privacy*, Volume 11, № 2 (2017): 25–37. PAYWALL

Şen, Fatih, Rolf Wigand, Nitin Agarwal, Serpil Tokdemir, and Rafal Kasprzyk,

'Focal Structures Analysis: Identifying Influential Sets of Individuals in a Social Network', *Social Network Analysis and Mining* Volume 6, № 1 (8 April 2016): 17. PAYWALL

Shu, Kai, Amy Sliva, Suhang Wang, Jiliang Tang, and Huan Liu. 'Fake News Detection on Social Media: A Data Mining Perspective', SIGKDD Explorations, 19(1):22-36, June, 2017.

Snopes.com, 'FACT CHECK: No More Child Support After 2017?', Kim LaCapria (fact-checker), 20 September 2017, updated 20 April 2018.

Snopes.com, 'FALSE: Comet Ping Pong Pizzeria Home to Child Abuse Ring Led by Hillary Clinton', Kim LaCapria (fact-checker), 21 November 2016, updated 4 December 2016.

Spence, Patric R., Kenneth Lachlan, Timothy Sellnow, Robert G. Rice, and Henry Seeger, 'That Is So Gross and I Have to Post About It: Exemplification Effects and User Comments on a News Story', *Southern Communication Journal*, Volume 82, № 1 (2017): 27–37. PAYWALL

Stats: A Look at Activity Across WordPress.com, provides a visualisation of current posts, comments, and likes, 8 December 2006. [last accessed 4 June 2018]

Subramanian, Samanth, 'The Macedonian Teens Who Mastered Fake News: Inside the Macedonian Fake News Complex' *Wired*, 15 February 2017.

TouchGraph, SEO Keyword Graph Visualisation. [Accessed 4 June 2018]

Vosoughi, Soroush, Deb Roy, and Sinan Aral, 'The Spread of True and False News Online', *Science* 359, № 6380 (2018): 1146–51.

Zillmann, Dolf, 'Exemplification Theory of Media Influence' Media Effects: Advances in Theory and Research 2 (2002): 19–41.