Blog Farm Detection Using Social Network Analysis and Social Cyber Forensics Informed Methodologies

Tuja Khaund[1], Billy Spann[1], Richard Young[1], Samer Al-Khateeb[2] and Nitin Agarwal[1]

[1] University of Arkansas at Little Rock, Little Rock AR 72204, USA
[2] Creighton University, Omaha NE 68178, USA
{txkhaund, bxspann, rbyoung, nxagarwal}@ualr.edu, sameral-khateeb1@creighton.edu

Abstract. Blogs play a vital role on the Internet for retrieval of real time information, a place for users to gain insights of events around them and also find communities that share similar interests. However, over the years, blogs have been in-fused with spam content and also have fallen prey to link farming. Blog search engines are plagued with fake content that may or may not be automatically generated with the motive of increasing the search ranking and thereby affecting crowd manipulation. In this paper, we monitor 25 blog sites obtained from the Twitter timelines of around 50 Ukraine parliament members and detect relevant blogs based on tracking codes. Our results show preliminary observations of the formation of a blog farm or a network of blog bots. We con-ducted a ping test on all the suspicious blogs and we observed patterns in their activity levels that suggest a presence of a master controller for the blog farms.

Keywords: Blog farms, Social Cyber Forensics, Social Network Analysis, Web Spam, Link farming, Crowd manipulation, Algorithmic bias, Information tactics.

1 Introduction

Social media platforms such as Facebook, Twitter, etc. have become a valuable resource for marketing, public relations etc., over the years. While there is still exchange of information that are relevant to users, spammers have taken over social media with their content spreading fake news and creating havoc on the web. Spam blogs and their presence on the web have immensely devalued blog search results and also have wasted plenty of network resources. Spammers also create fake blogs or submit spam comments/messages to host link farms [1]. Link farming on the Web refers to the exchange of reciprocal links by a website with other sites to improve ranking by search engines. Link farming previously referred to a form of spamming on search engine indexes that connected all of a webpage's hyperlinks to every other page in a group. Today, it's grown to include many graph-based applications within millions of nodes and billions of edges [2]. As noted in the literature, link farming can influence search rankings of a website. Malicious individuals or groups deploy link

farming to create a perception of a virtual crowd, making the content go viral, and hence effect crowd manipulation.

During NATO's Anakonda exercise in 2016, RT.com published an op-ed [3] that mocked Russia's involvement in the destruction of the Swedish TV Mast. The original article was published by Reuters accusing Russia [4] for vandalizing two telecommunications masts in Sweden. RT.com ridiculed the Reuters' article as a counter-attack to Sweden's effort to demonize Russia. RT.com posted their story on Twitter (Fig. 1, left) and several bots intensified the amplification (Fig. 1, right) to delegitimize the original article by Reuters. This amplification was so intense that even the Google's search results were affected. When Reuters published their article, Google would return their article as the top hit. However, after RT.com posted their article and it was intensively amplified via Twitter bots, Google showed RT.com's article as the top hit and Reuters' article was not even in the top few results. This and other similar accounts motivate the need for studying how link farming can help websites influence search engines and increase their ranks.

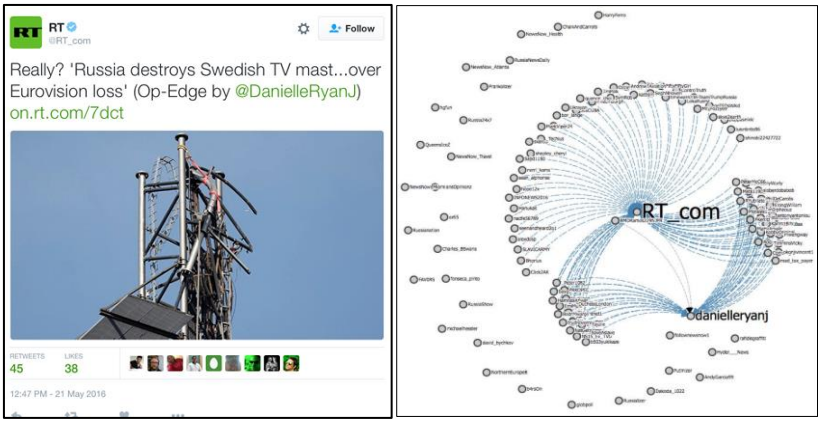

Fig. 1. RT.com's tweet posting the RT.com article on the Twitter to amplify its dissemination in order to delegitimize Reuters' article (left). Twitter network of the mentions of the RT.com's Swedish TV mast story and its author (@danielleryanj) (right)

Google Analytics tracking codes monitor the activity of a website and provide insights about visitors of the website. The Analytics tracking code may be added directly to the HTML of each page on a website, or indirectly using a tag management system such as Google Tag Manager. Domain Squatting, also known as Cybersquatting, is a process of buying a domain name with a bad intention to earn huge profits out of its direct or indirect usage in any way [5].

In this paper, we use social network analysis and social cyber forensics to detect blog farms on the web. This paper analyses a set of known blog sites obtained from Twitter accounts of the members of the Parliament in Ukraine and uses google analytics tracking codes to search for other relevant blogs. Social cyber forensics revealed blogs with unusual names connected to a single tracker code leading to a network

blog bots or blog farm. We also conducted a ping test on all the new blogs to check their status and behavior.

The rest of the paper is organized as follows. In the next section, we explore related work and examine the phenomena of spam blogs. Then we discuss our methodology where we use social cyber forensics to detect usual blogs based on tracking codes and followed by our results in sections 3 and 4. Finally, we present our discussion and future work in section 5.

2 Literature Review

Researchers have contributed towards detecting spammers in the blogosphere, but the majority of the work has been conducted on web spam detection. Gyongyi et al. [6] studied link farms on the Web, which are groups of interconnected web pages which attempt to boost the rankings of particular web pages. Specifically, they investigated how multiple web pages can be interconnected to optimize rankings.

Authors in [7] discuss that the basis of link farming on Twitter stems from how the Twitter search algorithm displays tweets. Content as well as the number of influential users that have posted a given tweet affects how highly a tweet is ranked when searched. Content spammers on Twitter aim to first acquire a large following to increase their perceptible "influence" within the twitter algorithm. This is done so that their tweets may reach a larger audience to promote their content spam. Analysis of the links formed between spammer accounts and the linked users show that most of the links are from legitimate, highly influential users. These users inadvertently resort to link farming by following back anyone who connects to them to increasing their influence. This behavior is mirrored and exploited by spam accounts. One method to discourage link farming could be a ranking system in which users would be punished for following spammers. This would have a negligent effect on real users, but a severe loss of influence amongst spammers [8].

Spam blogging has been considered as a special case of web spam pages [9, 10]. In order to combat spam blogs, the authors suggest using a set of content and link features and compare features in terms of classification performance by SVM classifier. In their work, each blog is treated as a single and static web page. Blogs have unique features. Unlike web spam where the content is usually static, a spam blog needs to have fresh content in order to continuously drive traffic and often the content is generated by an automated framework [11].

In a recent case study, the Atlantic Council's Digital Forensic Research Lab (DFRL) identified a network of at least 35,000 accounts, quite possibly over 75,000, whose purpose seems to be to steer users towards a cluster of pornographic chat sites registered in Russia. To mask their status as automated accounts, they quote snippets of text, many taken at random from Jane Austen's "Sense and Sensibility." [12] The network of dating and sex sites behind it appears to be similarly sprawling, with over a dozen sites all leading back to the central node. It was not certain if the botnet was primarily designed to make money from likes, retweets and follows, or by steering users to the pornography sites. It may even have served as a marketing A/B test, to

see whether more money can be made from Twitter-based bot activity, or off-site pornography. Either way, DFRL identified the botnet to be of commercial purpose, and detected cluster of sites on a large scale.

3 Methodology

While studying Ukrainian parliamentary discourse we identified a set of 25 blog sites from Twitter timelines of the parliament members of Ukraine. Using Maltego [13], a cyber-forensics tool, we first extracted tracking codes linked to these blogs. For the set of 25 blogs, only 17 blogs contained a tracking code. We selected those 17 blogs and used Maltego to detect other blogs that shared the same tracking code and identified several such instances. The blog sites may or may not have had a tracking code when the analysis was performed so we proceeded with the ones that had a tracking code. Several blog sites containing arbitrary numbers and characters were detected. These arbitrary blogs also appeared in alphabetical order with similar starting characters as part of the name, further suggesting they were part of a blog farm. The source of these blog sites was taken down after sometime, which led to a reverse track of its tracking code to retrieve the lost data. ORA [14] was used to visualize the network of blogs in which two blogs are connected if they share a tracker code. We also ran a ping test on the set of newly identified blog sites, once a day, starting from April 23, 2018 to monitor their activity. We present the analysis up to May 30, 2018.

4 Results

Maltego analysis identified a total of 141 new blogs from the initial set of 25 blogs. One popular Ukrainian news site named '*zolochiv.net'* shared a tracking code with another website named '*kinobot.ru*' that shared a tracking code with 100 new blog sites that had arbitrary characters and numbers (e.g. *000y.ru, 0058.ru, 000webhost.ru, 00000000000.ru, 00-00.net, 00300.ru, 000-shop.ru* to list a few). Figure 2 shows a snippet of the Maltego analysis of blog search based on tracking codes.

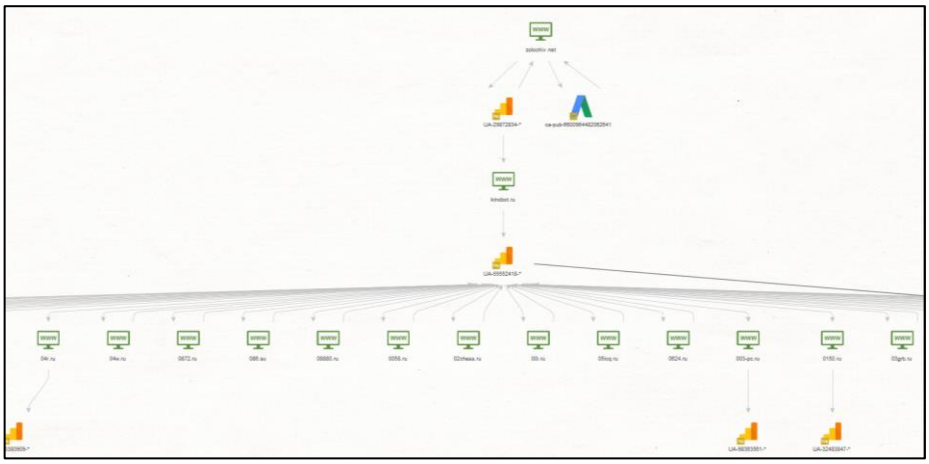

Fig. 2. Unusual Blogs identified by Maltego sharing the same tracking code.

Using ORA, we created the network of blogs connected by means of tracking codes (see Fig. 3.). The nodes are blogs and the edges are tracking codes. The nodes in aqua are the source blogs, blue nodes are the newly identified blogs and red nodes are the blogs that were newly identified that had arbitrary numbers or characters. This is the initial stage of the blog analysis where we first encountered these unusual blogs (the red nodes).

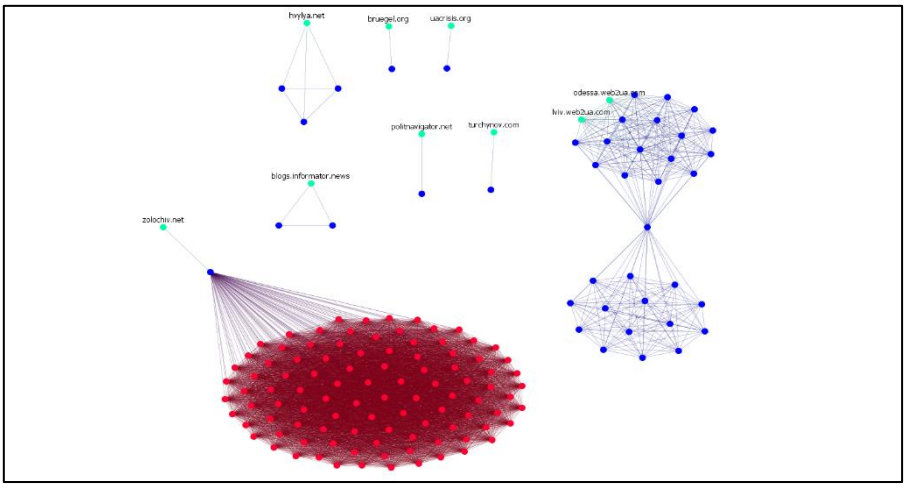

Fig. 3. Blog-Blog network sharing a tracking code.

As mentioned earlier, a blog farm of 100 such blogs were identified originating from a blog, '*kinobot.ru*'. Based on a manual check, most of these did not have any content or were domains for sale. Social Cyber Forensics was conducted again, a week later, on the same set of 25 blog sites which revealed that '*kinobot.ru*' was re-

moved and therefore, all the blog sites that were linked to it disappeared. This indicates that the blog farm was created with the intention of content pushing and amplification. The source blog may have been deactivated but its minions floated in the internet space. In order to retrieve the lost blog, we reverse tracked the tracking code linked to '*kinobot.ru*' and identified new blogs originating from the same tracker codes, however, the majority of these blogs were either empty or domains for sale. All but one of the blogs had the (***.ru***) top-level domain, and thus most of these blogs were registered in Russia. However, the cyber forensic analysis revealed that these arbitrary blogs were also registered in Japan, Netherlands, St. Petersburg, Russia, Germany, and the United States (see Fig. 4.).

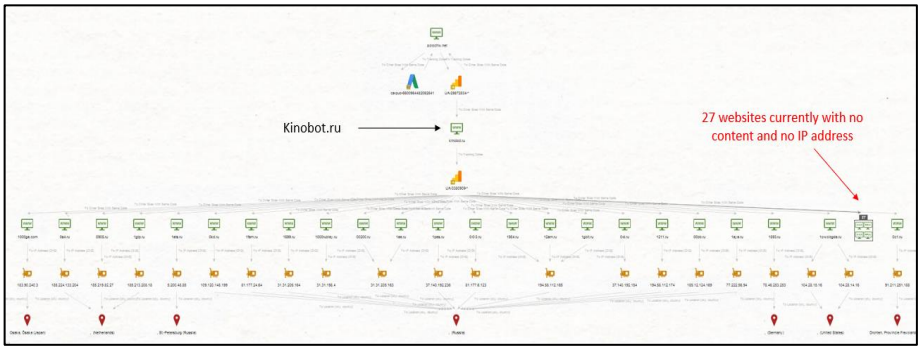

Fig. 4. Kinobot network showing multiple arbitrary blogs using same tracking code as Kinobot

We analyzed the rightmost cluster on Figure 3 and observed that one cluster consisted of websites with Ukrainian (***.ua***) domain and the other cluster consisted of websites with Russian (***.ru***) domain. The network in Figure 5 shows the nodes as the blog sites and the edges as the tracker codes shared between the nodes. The nodes are colored based on their PageRank centrality where red indicates the highest number of inlinks to the blog site and green indicates the lowest. This polarized behavior raised suspicions on the existence of blog farms that originated from two seed blogs, namely, *lviv.web2ua.com* and *odessa.web2ua.com*, as they were part of the *.ua* blog cluster and one of their derived blogs '*anapamoredom.ru*' that generated the *.ru* blog cluster (see Fig. 6.). The blog site (*anapamoredom.ru*) bridging these two clusters is simply a land infrastructure website owned by 'Field Sukko LLC' operating in cities such as Anapa and Sochi.

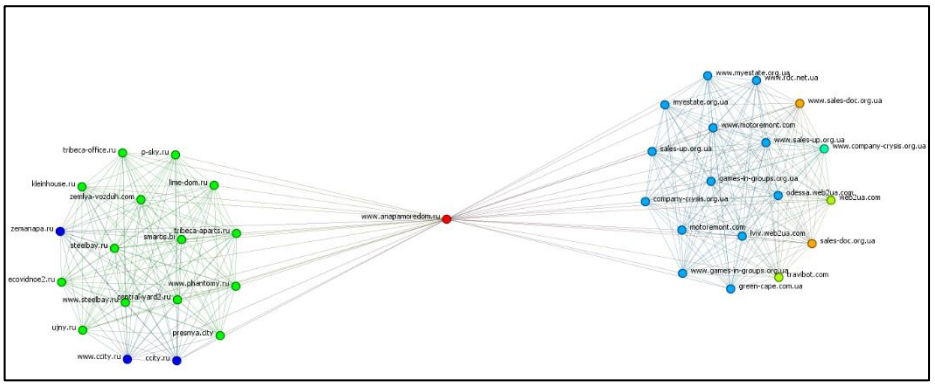

Fig. 5. Blogs indicating polarization of domains originating from seed blogs 'lviv.web2ua.com' and 'odessa.web2ua.com'

Fig. 6. Maltego cyber forensic representation of polarization of domains

Further analyzing the '*anapamoredom.ru*' bridge site (see Fig. 6.), we determined the IP address, and then used Maltego to run a cyber-forensic analysis to get several websites that mention '*anapamoredom.ru*' somewhere within those sites. We randomly selected one of the websites, *barrit.ru.pandastats.net*, and examined the content. This website appears to show web-traffic statistics for IP addresses, and one of the interesting statistics it showed was interest over time for the IP address of our bridge site, *anapamoredom.ru.* The interest over time statistic is a function of the Google Trends metric that shows how frequently a given search term is searched. Using our example, and referring to the graph in Figure 7, we see that the search term *barrit* had a spike in traffic in September of 2015. This is the same time period where a ceasefire agreement was broken in Ukraine by Russian-backed separatist forces [15]. If the *anapamoredom.ru* blog was part of a larger blog farm as shown in the social cyber forensics analyses, then the spike in traffic during the time of the crisis could indicate that the website was attempting to amplify content. Additional research, beyond the scope of this paper, is needed to determine if this content spike is related to the ceasefire crisis or a larger collective operation to amplify content.

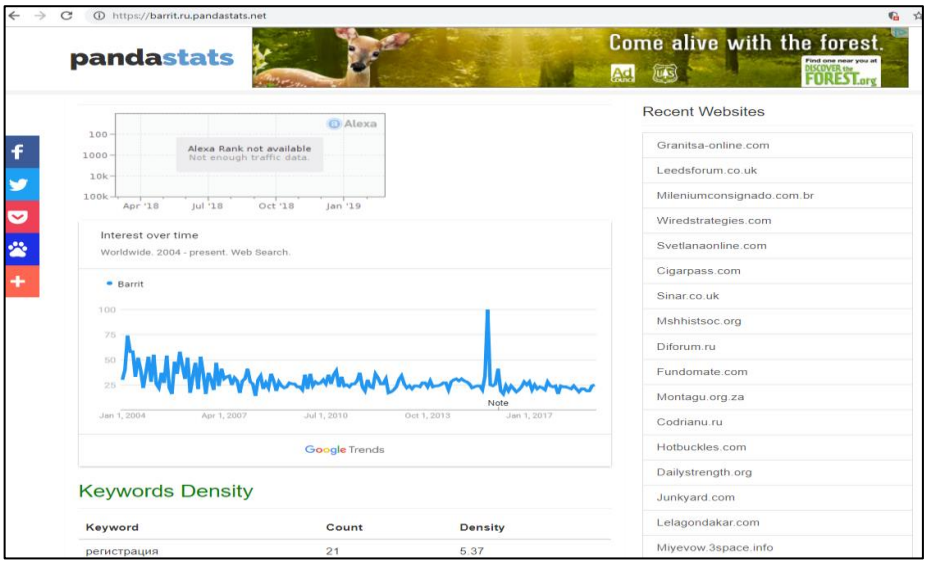

Fig. 7. Content of website *barrit.ru.pandastats.net* revealing spike in traffic in Sep.2015

To dig deeper into the activity trends of the website, we analyzed the ping responses (such as *active*, *host-not-found*, and *timed-out*) for the blog sites. The scatterplot in Figure 8 allowed us to quickly visualize and identify any hosts with fluctuations in their activity. We then manually filtered for the hosts with interesting behavior and generated the scatterplot seen above. The green points indicate *active*, red points indicate *host-not-found* or inactive and blue points indicate *timed-out*.

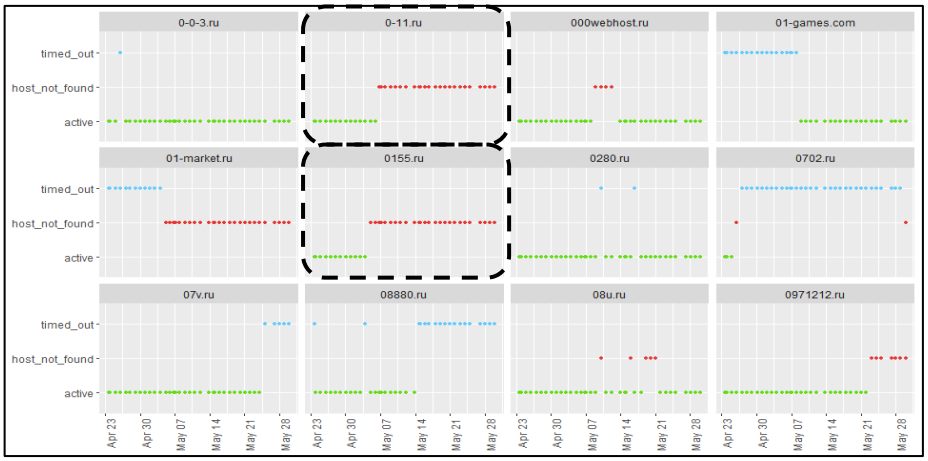

Fig. 8. Ping test results for a sample of the blogs that kept fluctuating.

We observed highly similar behavior in two blog sites, viz. *0-11.ru* and *0155.ru* (marked with dashed rectangles in Fig. 8), i.e., ping responses over the course of one

month were identical for these sites, which indicate that these two sites might have been controlled by the same program or person. Since most of these blogs sites do not have any content yet, we assume it to be *sleeper cells* that stay inactive or under the radar until they are activated by their controller.

5 Discussion and Conclusions

In this paper, we monitor blogs and detect relevant blogs based on tracking codes. Our results show preliminary observations of the formation of a blog farm or a network of blog bots. We conducted a ping test on all the suspicious blogs and we observed fluctuations in their response. The Majority of these blogs did not have any content yet but we believe these arbitrary blogs were created to post spam content that are targeted to specific events. There were plenty of domains for sale which could be a possible case of Cybersquatting. Majority of these blogs' IP addresses were from Russia, Ukraine and the United States. Also, many of these blog sites were detected from a single tracking code but the relation is one way. This is because there wasn't any content available and hence the HTML page did not have JavaScript code embedded to it even though these sites were registered to be tracked.

Unlike spam content that can be caught via existing anti-spam techniques, link-farming fraudsters can easily avoid content-based detection, for example, in Twitter's "who-follows-whom" graph, fraudsters are aid to make accounts seem more legitimate by giving them additional followers, called zombies. Twitter zombie followers don't have to post suspicious content, they just distort the graph structure. Thus, the problem of combating link farming is rather challenging[2].

For our future work, we would like to investigate deeper into these link farms and cross check on other platforms such as Twitter. We are also monitoring their content status and with new content, we will be able to study their motives better. For now, our experiment showed potential blog farms that are currently hiding under the radar until activated.

Acknowledgements

This research is funded in part by the U.S. National Science Foundation (IIS-1636933, ACI-1429160, and IIS-1110868), U.S. 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), U.S. Air Force Research Lab, U.S. Army Research Office (W911NF-16-1-0189), U.S. Defense Advanced Research Projects Agency (W31P4Q-17-C-0059), Jerry L. Maulden/Entergy Endowment at the University of Arkansas at Little Rock and the Arkansas Research Alliance. 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 organizations. The researchers gratefully acknowledge the support.

References

1. Agarwal, N., Yiliyasi, Y.: INFORMATION QUALITY CHALLENGES IN SOCIAL MEDIA. 16.
2. Jiang, M., Cui, P., Faloutsos, C.: Suspicious Behavior Detection: Current Trends and Future Directions. IEEE Intell. Syst. 31, 31–39 (2016).
3. Really? “Russia destroys Swedish TV mast ... over Eurovision loss,” https://www.rt.com/op-ed/343901-putin-destroys-swedish-mast/.
4. Sabotage of telecoms masts reignite Swedish security fears, https://www.reuters.com/article/us-sweden-masts/sabotage-of-telecoms-masts-reignite-swedish-security-fears-idUSKCN0Y921R, (2016).
5. Oketunji, F.: Domain Squatting: Everything You Need To Know, https://techcabal.com/2014/10/13/domain-squatting-everything-need-know/, (2014).
6. Gyongyi, Z., Garcia-Molina, H., Univ, S., Pedersen, J.: Combating Web Spam with TrustRank. 12.
7. Ghosh, S., Viswanath, B., Kooti, F., Sharma, N.K., Korlam, G., Benevenuto, F., Ganguly, N., Gummadi, K.P.: Understanding and combating link farming in the twitter social network. Presented at the (2012).
8. Agarwal, N., Liu, H., Tang, L., Yu, P.S.: Modeling blogger influence in a community. Soc. Netw. Anal. Min. 2, 139–162 (2012).
9. Kolari, P.: Detecting Spam Blogs: A Machine Learning Approach. 6 (2000).
10. Kolari, P.: SVMs for the Blogosphere: Blog Identification and Splog Detection. 8.
11. Lin, Y.-R., Sundaram, H., Chi, Y., Tatemura, J., Tseng, B.L.: Splog detection using self-similarity analysis on blog temporal dynamics. Presented at the (2007).
12. Nimmo, B.: #BotSpot: Sex And Sensibility, https://medium.com/dfrlab/botspot-sex-and-sensibility-dc1d4a72a92e, (2018).
13. Paterva Home, https://www.paterva.com/web7/.
14. Projects - *ORA-LITE | CASOS, http://www.casos.cs.cmu.edu/projects/ora/.
15. Gibbons-Neff, T.: Three-day-old ceasefire in Ukraine broken as fighting resumes in some areas.