Characterizing Search-Engine Traffic to Internet Research Agency Web Properties

Alexander Spangher[1], Gireeja Ranade[2,3], Besmira Nushi[3], Adam Fourney[3], and Eric Horvitz[3]
[1]Information Sciences Institute, University of Southern California
[2]University of California, Berkeley
[3]Microsoft Research

ABSTRACT

The Russia-based Internet Research Agency (IRA) carried out a broad information campaign in the U.S. before and after the 2016 presidential election. The organization created an expansive set of internet properties: web domains, Facebook pages, and Twitter bots, which received traffic via purchased Facebook ads, tweets, and search engines indexing their domains. In this paper, we focus on IRA activities that received exposure through search engines, by joining data from Facebook and Twitter with logs from the Internet Explorer 11 and Edge browsers and the Bing.com search engine.

We find that a substantial volume of Russian content was apolitical and emotionally-neutral in nature. Our observations demonstrate that such content gave IRA web-properties considerable exposure through search-engines and brought readers to websites hosting inflammatory content and engagement hooks. Our findings show that, like social media, web search also directed traffic to IRA generated web content, and the resultant traffic patterns are distinct from those of social media.

ACM Reference Format:
Alexander Spangher[1], Gireeja Ranade[2,3], Besmira Nushi[3], Adam Fourney[3], and Eric Horvitz[3]. 2020. Characterizing Search-Engine Traffic to Internet Research Agency Web Properties. In *Proceedings of The Web Conference 2020 (WWW '20), April 20–24, 2020, Taipei, Taiwan.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3366423.3380290

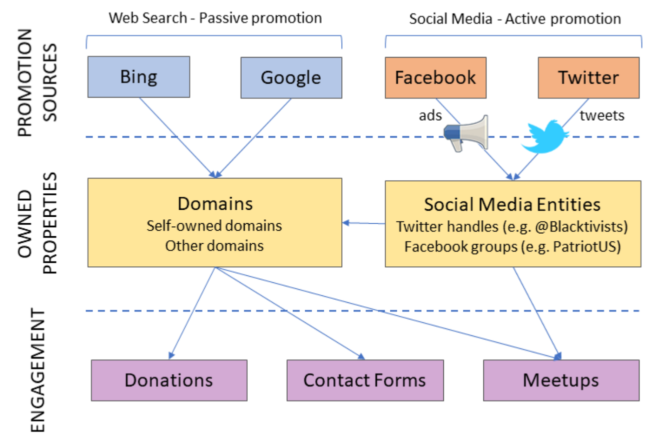

Figure 1: IRA campaign structure. Illustration of the structure of IRA sponsored content on the web. Content spread via a combination of paid promotions (Facebook ads), unpaid promotions (tweets and Facebook posts), and search referrals (organic search and recommendations). These pathways pointed to a combination of pages, some of which were created by the IRA (e.g. blackmattersus.com).

1 INTRODUCTION

The Internet Research Agency (IRA), a Russia-based company focused on media and information propagation [29], was found to have spent at least 100, 000.36 USD from June 8, 2015 to August 1, 2017 [36], to disseminate information to the U.S. public via Facebook advertising. They fielded thousands of Facebook advertisements and sent millions of Tweets, promoting hundreds of web domains and Facebook groups that spanned the political spectrum [30, 42, 43]. According to an indictment made against the IRA by the U.S. Special Counsel's Office on February 16, 2018, the purpose of IRA's activities was to "sow discord in the U.S. political system, including the 2016 U.S. presidential election" [35].

In the process of investigating Russian election interference, the U.S. House Intelligence Committee released IRA-linked Facebook advertisements and Twitter accounts into the public domain [42, 43]. As stated in the release:

> Russia exploited real vulnerabilities that exist across online platforms and we must identify, expose, and defend ourselves against similar covert influence operations in the future.

Extensive research has focused on the IRA's use of social media platforms, primarily their production of inflammatory content and purchase of advertisements [4, 12, 13]. However, the IRA's activities were not focused only on social media. According to the February 2018 indictment by the Special Counsel's office: the IRA was *"organized into departments, including: a graphics department; a data analysis department; a search-engine optimization (SEO) department..."* [35]. A major gap in our understanding of the IRA's activities centers around how content was accessed by search-engines.

In this paper we focus on identifying tactics targeting web search platforms, utilizing datasets heretofore unconsidered for this purpose. We find tactics that contrast with their tactics on social media: *the IRA produced a volume of apolitical, emotionally-neutral and news-driven content which performed well on search engines and were sometimes featured in company promotional channels.* Although we

WWW '20, April 20–24, 2020, Taipei, Taiwan

ACM ISBN 978-1-4503-7023-3/20/04.
https://doi.org/10.1145/3366423.3380290

did not find evidence of paid advertisements on search, our findings indicate that search-engines were an important part of the IRA's campaign.

Our analysis considers 212 days of IRA activity in 2017 (January 1 to August 1, 2017). We merge publicly-available Facebook and Twitter datasets with traffic datasets: logs from Internet Explorer 11 and the Bing search engine, thus fusing data from three major internet companies to provide a broader lens on the scope and influence of the IRA generated content (Fig. 1). We show in case studies that apolitical content and local news produced by related IRA-properties reached users through search, and played a role in bringing traffic to IRA domains. We further examine the flow of traffic once users arrived on these domains.

Our specific findings are as follows:

(1) *IRA Content:* The IRA produced and amplified a diverse array of left-leaning, right-leaning and apolitical content, which ranged from emotionally charged to neutral (Fig. 3). IRA content drew traffic mainly from social media (80.7 % of traffic) and search domains (11.0 % of traffic).
(2) *IRA Strategies*: We identify a strategy of employing apolitical content to garner credibility and engagement; such content received consistent traffic via *organic* searches (both *informational* and *navigational*), but also received large spikes in traffic via indirect promotional channels, which lead users to search-engine result pages (SERP) powered by *predefined* queries[1]. For example, one IRA article about Black Female Computer Scientists[2] appeared among the search results returned for the query "women scientists" that was promoted across device lock-screens (Fig. 7 and Section 6). We also examine a tactic employed via Twitter that performed well on search: rapid tweeting about evolving local news stories (Fig. 8). Such tweets were published before traditional local news published stories on the same topic appeared on SERPs.
(3) *IRA Traffic Outcomes:* Despite the large volume of content generated by the IRA, on average roughly 1 in 40,000 internet users in our dataset clicked on an IRA-property (Tweet, Facebook Group, URL) on a given day. Only half of these users, or 1 in 80,000, clicked on more than one IRA-property in a day and only 0.02% of these users clicked on both IRA Facebook and IRA Twitter content. This warrants further research that could take into account what we could not observe, such as traffic before and during the election, broader traffic from multiple browsers and search engines, traffic from mobile applications, long-term user behavioral changes, and real-world actions.

The remainder of the work is structured as follows: we contextualize our work with prior research in this space (Section 2) and we describe our datasets and methodology (Section 3). Next, we contextualize the importance of our analysis focused on search by describing the high-level structure of the IRA web properties and promotions (Section 4), and show traffic outcomes (Section 5). Our observations are that search traffic drove a significant minority of traffic to IRA properties, and that neutral and apolitical content formed a significant portion of the content published by the IRA.

We then present results and case-studies showing tactics employed to garner traffic through search (Section 6) and show the traffic results of these tactics. Finally, we discuss the implications and limitations of this research, as well as future directions.

2 RELATED WORK

The IRA Campaign and Tactics. Our work fits into a larger body of journalistic, governmental and academic research focused on understanding the IRA's campaign by using traffic and search datasets to confirm and expand previous observations.

The IRA first came to light after the *New York Times* reported on its existence in June 2015 [9]. Multiple journalistic efforts have since given illustrative examples of the IRA's content [32, 44], and the U.S. Special Counsel's February 2018 indictment which further exposed the IRA's organizational structure [35]. In May 2018, the United States House of Representatives Permanent Select Committee on Intelligence released IRA Facebook advertisements and IRA Twitter handles into the public domain [42, 43]. These were soon further expanded to include IRA tweets [30].

Using these datasets, academic researchers and journalists examined the structure of the IRA's social media campaign, the IRA's retweet network [1], characteristics of IRA-promoted Twitter handles [30] and how these handles differed from organic Twitter users [27]. Researchers also examined Facebook ad effectiveness [13] and divisiveness [41]. Additionally, separate work by our group focused on further analyzing the point-of-origin of IRA activity [6].

Another dataset compiled in December 2018 by the United States Senate Select Committee on Intelligence, as yet unreleased, supplemented the House Intelligence findings and included content from Facebook, Twitter, YouTube, and Instagram. Researchers analyzed it to uncover additional IRA campaign structure [12, 26], showing different tactics used on different platforms to attract audiences. Finally, in April 2019, the U.S. Department of Justice released the "Report On The Investigation Into Russian Interference In The 2016 Presidential Election", or the Mueller Report [36], providing a detailed forensic analysis on IRA activities throughout 2016.

Collectively, a nuanced portrait of the IRA and its tactics has emerged.The IRA's campaign builds off a long history – disinformation was used by state-governments as early as 1964 to influence popular opinion [15]. Classical disinformation efforts relied several key tactics: (1) Soviet KGB intelligence services promoted disinformation claims over long periods of time, through multiple sources. [46]. (2) These claims often first appeared in stories planted in local news outlets and tied to local issues, making them seem "boutique" to (future) global audiences [31]. (3) Alternatively, claims were spread through seemingly apolitical organizations, like the World Peace Congress [5] and the Russian Orthodox Church [37]. (4) Claims were then cross-cited and expanded through different journalistic and academic channels, giving credibility and global reach [15].

[1] An *organic* search is a user-typed query typed into a query bar. An *informational* organic search is a general information-seeking query. A *navigational* organic search, in contrast, is typed into the query bar as well, but is intended to bring the user to a specific web property (ex. "nytimes" or "cnn"). A *predefined* query, further defined later, is an auto-generated query (i.e. not typed by the user) that a user navigates to by clicking on another Microsoft property, like a device lock-screen or news page.

[2] "Black Female Computer Scientists. How Many Do You Know?" published at https://blackmattersus.com/31081-black-female-computer-scientists-how-many-do-you-know/. The promotional query leading to this page was "women scientists".

Many of these same tactics, as noted in Section 1, are adapted by the IRA for their online campaign: the use of apolitical content, local news content, and a cross-platform approach. By layering additional traffic data from internet browsing logs and search engines, our work is the first known work to show the traffic-outcomes of such tactics. Our user-behavior insights confirm and expand observations made previously and help paint a more detailed picture of the IRA's campaign.

Disinformation, Search, and Social media. The data made available during Congressional inquiry into the IRA provides a unique opportunity to study how a single agent can carry out active measures in the modern age through social media services. It fits within a larger body of work examining the impact of social media and search engines on information consumption.

Social networks have facilitated faster and more targeted user communication, which has created new and amplified effects [2]. From a political perspective, such effects can often be positive and increase political participation and awareness [25]. However, they can also create drastic political division [48] and spread misinformation [19, 33]. Political division has been investigated using the notion of "information bubbles" [40] and "echo chambers" [3, 20]. While some of these behaviors can occur organically on social platforms, research has shown that malicious agents can work to increase division [3, 10, 39, 47].

The role of search engines in information consumption has been widely studied: research into ranking effects on search-engines has identified the Search Engine Manipulation Effect (SEME), where variations in ranking results in shifts in user opinion [14, 17]. Researchers have examined success-rates of actors spreading misinformation on search engines in the 2016 election [34], finding that misinformation was spread without proper context on major search engines. Our work continues this vein of inquiry and utilizes search-engine datasets to highlight not just the success of misinformation on platforms but the tactics that actors might have employed to gain reach.

Various approaches have been proposed to identify disinformation using linguistic features, semantic analysis of the content, and network topology properties [8, 11, 45]. Other approaches have been proposed to help internet users guard against disinformation's harmful effects. Proposals have been made to foster diverse information consumption either via tools for visualizing Twitter feeds [22] or content exploration strategies that go beyond the personal network bubble [40]. Other research have shown that when search-engine results are paired with warnings or background information, the SEME of a highly ranked piece of disinformation can be mitigated [18]. We find this direction particularly exciting and we hope that this current work, in its quantitative approach to analyzing the IRA's social media campaign, can inspire further approaches for combating disinformation online.

3 DATA AND METHODOLOGY

We have harnessed data from three large internet companies – Facebook, Twitter and Microsoft – to provide a perspective on the strategies employed by the IRA from January 2017 to August 2017. This work also leverages additional open-source and third-party data sources. We describe each data source in turn.

3.1 Primary datasets

Facebook ads: On May 10, 2018, the House of Representatives Permanent Select Committee on Intelligence released 3,571 Facebook advertisements reported by Facebook as paid for by IRA-linked entities as PDF documents[3] [42]. We performed optical character recognition (OCR) and template-based information extraction, and this process yielded 3,296 ads with actionable data.

Of the 275 ads without actionable data, 30 of these ads are blank or contain redacted data. We communicated with researchers from Facebook to verify these numbers. Researchers confirmed that redactions were performed to protect genuine Facebook users who were unwittingly promoted by the IRA. We performed a manual check on the remaining 245 ads and confirmed that, because multiple ads may promote a common landing page, we were missing an insignificant number of landing page URLs.

IRA-linked Tweets: Alongside the Facebook ads data, the House Intelligence Committee also released a dataset on June 8th consisting of 3,841 Twitter account handles believed to be associated with the IRA [43]. Researchers at Clemson University scraped over 2.9 million tweets from these handles and performed language and topic classification [30]. We leverage their work and focus our analysis on English-language tweets occurring from January 1st to August 1st, 2017 from the following categories they assign: "LeftTrolls", "RightTrolls", "Local" and "News" (See [30] for more details). This resulted in a total of $471,000$ tweets and 320 handles. We confirmed these numbers and figures with Congressional staff.

Additionally, many of the tweets link to external domains. We augment the data by resolving Twitter's link-shorteners using historic instrumentation logs from Microsoft web browsers to observe where users clicking on these links were directed. These logs are described next.

Browsing data: We consider 212 days (January 1 to August 1, 2017) of instrumentation data collected by Microsoft Edge and Internet Explorer 11 desktop web browsers. [4] Data includes anonymized time-stamped records of page visits. Records are assembled into sessions of browsing activity: a "session" is defined as a sequence of page visits such that consecutive visits are less than 30 min. apart.

Web search data: Finally, we also analyze logs from the Bing search engine. This data reveals URLs in Bing search results that were clicked on by users (i.e. *clicks*) as well as links presented as search results that were not ultimately clicked by users (i.e. *impressions*) – something that the browsing instrumentation data cannot provide. This dataset considers the same 212 days of data as the browser instrumentation logs. Both sets of traffic data include users' approximate geo-locations.

3.2 Joins of primary datasets

We join the primary datasets to recover the scope of the IRA campaign. As shown in Figure 2, the joins we performed are represented by intersections of shaded blocks (Internet Explorer/Edge clicks, Bing clicks and Bing Impressions) and squares (Facebook ads, Tweets, and landing pages).

[3] PDF documents for the individual ads included: text associated with the ad, the ad's start and end dates, targeting information, ad-spend, and the URL of the web property that the ad was promoting.

[4] Browser data is collected anonymously with user permission.

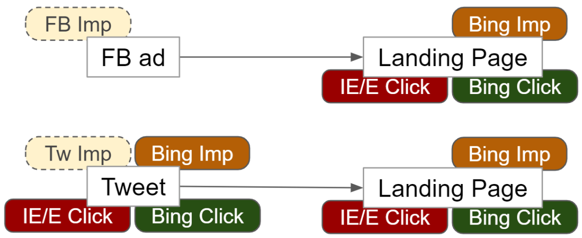

Figure 2: Data flowchart: The datasets available for our analysis were Internet Explorer/Edge clicks, Bing clicks and Bing impressions on Facebook-promoted landing pages (above), tweets, and tweet-promoted landing pages (below). We did not have access to impressions on Facebook ads, nor impressions on tweets (dashed lines). Rectangles signify content, rounded rectangles signify behavioral data, and arrows signify results when ads or tweets are clicked.

A key source of data – impressions of Facebook ads and tweets occurring on Facebook and Twitter – was not observable to us. Therefore, we associated IRA Facebook advertisements with Internet Explorer/Edge browsing data by identifying clicks from `facebook.com` to a URL promoted by an IRA ad[5]. An equivalent approach was followed for joining browsing data with Twitter account handles and tweets[6].

Bing search engine logs were joined to the social media data by matching URLs of search results to URLs of pages actively promoted by Facebook ads, and URLs of tweets sent from IRA-linked Twitter accounts (tweets are indexed and surfaced in Bing SERPs). In search logs, we detect when links are presented and when they are clicked.

Collectively, we observe hundreds of thousands of users and millions of clicks on IRA content, with comparable proportions across Facebook and Twitter datasets.

3.3 Crowdsourced labeling of primary data

We used crowdworkers to label content[7] to ascertain information about the IRA content's political leaning and emotional intensity. We ran this study for Facebook ads, tweets, and Bing search snippets generated for URLs on IRA-owned domains.

For Facebook ads, we rated all 1,032 ads that ran after January 1, 2017. For Twitter we rated two subsets: a random subset of 500 IRA-linked tweets, and the top 500 most-clicked tweets. For Search, we rated the 100 URLs with the most impressions (these 100 URLs captured 63% of search-result impressions to URLs in our datasets).

[5]After normalizing and reversing URL-transformations that Facebook applies to URLs (e.g. splicing /pg/ or /?/ from URLs). While such events are consistent with users clicking on ads, they are not sufficient to conclude that a particular ad was clicked because (a) a user may have arrived at a page from elsewhere in Facebook (e.g. from the newsfeed, or a notification) and (b) multiple ads promote a common URL.

[6]We searched for clicks with twitter.com and <accounthandle> as substrings, thus capturing clicks on tweets *and* account handles.

[7]We use Amazon Mechanical Turk. Workers were paid on average $12 per hour, above the national minimum wage at the time of publication. We limit for each crowdworker to 50 tasks each, and surveyed a total of 604 workers. We asked each worker their state of residence once (U.S.), in order to test how closely the crowd-worker population matches our observed browsing data. We find no significant geographical difference ($\chi^2 = .19$, $p = .999$) between the states our crowdworkers live in and the states our internet users live in. Tasks were completed with a median time of 73 seconds (Q1: 31 seconds, Q3: 491 seconds).

The task showed workers the text content of the promotion: the original Facebook ad, tweet text, or the snippet text shown in search results. Based on this information, the workers' task was to label the content with the most relevant option in each of the following categories: (1) *Political leaning*: {extreme-left, left, center, right, extreme-right, apolitical, local news}. (2) *Emotional intensity*: {neutral, low, medium, high, very high}[8]. We gave minimal information to define the *Political leaning* categories for workers[9] and no definitions for *Emotional intensity* – we sought to leverage workers apriori definitions of these labels, as this mirrored the experience of internet users in our dataset encountering IRA content.

Each item was judged by five different crowd workers. We performed quality control by statistically identifying and excluding answers from workers with high disagreement rates, (according to methodology by [28]), and performed a soft-assignment of these worker tags to label each item[10]. To calculate the number of clicks on a certain label, we used the soft assignment to divide the clicks across the assigned labels, and ultimately combined some categories to reduce sparsity.[11]

3.4 Secondary datasets

In addition to the primary datasets, we employ external secondary datasets. We use SimilarWeb Domain Categories [38], which provide domain-level labels to characterize the content of URLs in our datasets, and GDELT Global-Events data, which captures media publications over time, to characterize external news events [21].

3.5 Limitations

Much of the findings we present must be viewed with the following limitations. First, we use Twitter handles and Facebook ads identified by these companies: we rely on their methodologies for identifying suspicious activity as we lack the data and insights necessary to replicate their work. To mitigate this shortcoming, we successfully verified our results with representatives from Facebook, and we attempted to verify our results with representatives from Twitter. We also used the most recently updated lists of IRA Facebook ads and IRA Twitter handles released by Congress.[12]

Second, since our browsing instrumentation collects only URLs, our browser logs cannot directly observe interactions that occur within a page[13]. Additionally, the search engine we study surfaces URLs and tweets, not Facebook ads. Thus, we constrict our analysis

[8]The questions asked, specifically, were "*What are the politics of this piece of content?*" and "*What is the emotional intensity of this piece of content?*".

[9]For labels on the political spectrum, we defined them in relation to the political parties. For example, we defined "extreme-left" as "*appears to support policy more liberal than current Democratic Party stances*" and "right" as "*appears to support Republican Party/conservative policy issues*". For "local news", our definition was: "*primarily involves local political issues or issues that do not have national prominence*". For "apolitical", our definition was: "*Does not contain any political message*".

[10]For example, if a certain URL received 2 votes for extreme-right, 1 for left and 2 for extreme-left, then we counted 2/5, 1/5 and 2/5 clicks in each category respectively.

[11]A vote for either extreme-left or left was considered "left", extreme-right and right for "right", apolitical and local news for "apolitical", and high and very-high emotion for "high".

[12]An earlier list of Twitter handles released by Congress contained a small number of genuine, non-IRA U.S.-based users and did not identify an additional 1,103 suspected accounts (see: https://www.wired.com/story/how-americans-wound-up-on-twitters-list-of-russian-bots/).

[13]E.g., interactions enabled by AJAX or JavaScript are not observable to us, like scrolling down a Twitter or Facebook feed.

to (a) landing pages and domains identified in the Facebook-ad dataset and (b) tweets.

Third, this analysis cannot observe the long-term behavior of internet users online due to anonymization techniques applied in the data. Thus, our observations on cross-platform engagement (Section 5) need to be viewed with this important caveat.

Fourth, we do not generalize our findings beyond the population of internet users we studied[14]. Additionally, while the web-search dataset we use includes mobile traffic, the browsing instrumentation data is derived from desktop computers and excludes mobile devices. A complete analysis on these effects calls for data from potentially multiple browsers and search engines to better cover the online user base, as well as data from a broader range of mobile devices and applications.

Fifth, browser and query logs give us only a limited view into users' online actions (e.g. public posts and tweets). For example, we cannot be sure if a person shared news verbally, or via direct message. Direct message, in particular, was a prominent tool used by the IRA to reach out to audiences [35]

Finally, we do not incorporate insight into numerous confounders that may have existed in this complex system. Actions taken by non-IRA entities may have had effects on traffic: for example, the 2018 Mueller Report points to examples of supportive actions: "*in total, Trump Campaign affiliates* ***promoted dozens of tweets, posts, and other political content created by the IRA*** *(...)*" [36]. Additionally, countermeasures by third parties, like monitoring or active intervention by any administrative body – corporate or governmental – were also not considered.

4 STRUCTURE OF THE IRA PROMOTIONS

We first present an overview of the scope and structure of the IRA's campaign to illustrate it's breadth across different platforms.

4.1 Facebook Ads

We found that the Facebook advertisements paid for by the IRA promoted a smaller set of 350 distinct URLs. These URLs were distributed across 225 web properties. These properties include Facebook groups and profiles; Facebook event or meetup.com pages; news websites (including CNN.com); petitions (whitehouse.gov, change.org); four domains identified as controlled by the IRA.[15]

For this paper, we limit our analysis to internet traffic to landing pages promoted by the IRA's Facebook campaign that are suspected to be IRA-controlled[16]. However, the breadth of different content-types promoted in the IRA's Facebook campaign is significant (petitions, events, articles, Facebook groups), and we note that the reach of the IRA campaign went beyond what can be measured by simply analyzing online traffic patterns. For instance, one of the IRA's petitions[17] on change.org received 65,000 supporters, within the range of mid-to-high performing change.org campaigns [16]. Additionally, the IRA campaign received voluminous news coverage worldwide, often being mentioned in thousands of news articles a week over months (as shown in Figure 9).

4.2 Twitter

The $3,841$ Twitter handles released by [43] and 2.9 million tweets compiled by [30] were filtered to $471,000$ English-language tweets from 320 "LeftTrolls", "RightTrolls", "Local" and "News" accounts. Our analysis of links in these tweets revealed links to over $5,500$ domains, tweets, and other properties. Over 95% of tweets are retweets, or reposting of other users' tweets, and many of the tweets point to well-known websites, like nytimes.com.

As noted by [12], overlap in naming exists between the IRA's Facebook and Twitter campaigns: for instance, the Twitter campaign included a Twitter handle "blackmattersussoldier" and the Facebook campaign included a "BlackMattersUS" Facebook Group. However, we do not see the Facebook and Twitter campaigns sharing links or cross-promoting each other. We observed only four URLs that were promoted both by Facebook ads and tweets, all of them blackmattersus.com URLs with small volumes of traffic.

Only a small fraction of the landing pages promoted by the IRA's tweets were IRA-controlled (in contrast to the landing pages promoted by the Facebook ad campaign, which almost all IRA-controlled). Many of the links promoted by IRA tweets included major domains, like nytimes.com mentioned above, which were promoted by many actors external to the IRA. As a result, we focus on an analysis of the tweets themselves (for the subset outlined in Section 3), and not the domains they promoted. Thus, throughout this paper, when we refer to "IRA Twitter content", we are referring to IRA tweets and handles that were identified in the Twitter dataset.

4.3 Web Search

We find that many of the web properties promoted by the IRA in their Facebook ad campaign were indexed by search engines, and surfaced to users in response to web search queries. Many of the most-surfaced pages were suspected to be IRA-owned domains (with the notable exception of a guardian.com article, which we exclude from our analyses). Of the suspected IRA-owned domains we find that one domain in particular, blackmattersus.com, received significant traffic via the Bing search-engine (Fig. 5 and Fig. 7). However, a small amount of SERPs also contained IRA-controlled Facebook properties[18].

Additionally, we find that IRA tweets were exposed directly to users via the Bing search-engine. Search-engines typically surface tweets along with other forms of information, especially for fast-developing news-stories when tweets can give readers rapid and up-to-date information. In this manner, IRA tweets gained traffic through these channels.

We did not find evidence of paid advertising on the search-engine: we study this by examining search logs for any clicks on a paid-ad containing suspected IRA domains. A more comprehensive evaluation of purchasing records might yield further evidence.

[14] I.e., we do not account for behavioral factors that might lead a user to choose to use the platforms we studied compared with other platforms.

[15] We identified: blackmattersus.com, dudeers.com, black4black.com, donotshoot.us, see [12] for a more complete list.

[16] Due to our browser limitations, mentioned in Section 3, we could not track when users saw and clicked directly on IRA Facebook ads. Facebook ads are rendered via Javascript on facebook.com and are thus not tracked in our browser logs, as depicted in Figure 2

[17] https://www.change.org/p/barack-obama-u-s-house-of-representatives-u-s-senate-list-the-ku-klux-klan-as-an-official-terrorist-organization

[18] E.g. facebook.com/blackmattersus.mvmnt and facebook.com/events/535931469910916

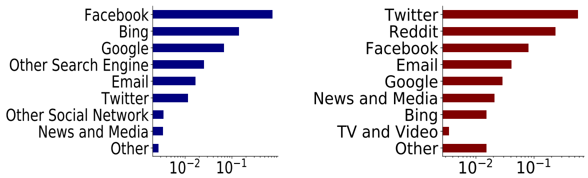

(a) Referral pathways for clicks on URLs included in the IRA Facebook ads. **(b) Referral pathways for clicks on tweets linked to IRA accounts.**

Figure 3: Top traffic channels to IRA properties. Top sources that brought users to IRA properties (Facebook-advertised URLs, Facebook groups and tweets) from January 1, 2017 to August 1, 2017.

Top-Clicked Facebook Groups	Share of IRA Traffic
blacktivists	0.208
godblessthesouth	0.199
blackmattersus.mvmnt	0.169
brownunitedfront	0.116
patriototus	0.077
Other	0.231
Top-Clicked Twitter Handles	**Share of IRA Traffic**
ten_gop	0.462
pamela_moore13	0.245
crystal1johnson	0.076
southlonestar	0.048
jenn_abrams	0.045
Other	0.123

Table 1: Top-clicked IRA properties (across all channels). The top five Facebook groups (top) and Twitter handles (bottom) with most traffic. These Facebook groups collectively account for 77% of traffic to IRA-promoted Facebook properties, while these Twitter handles account for 88% of traffic.

5 IRA OVERALL TRAFFIC PATTERNS

Having described the breadth and size of the IRA's campaign in Section 4, we now summarize the content of the campaign and provide an overview of the traffic to it. Our results in this section provide a motivation for the study of search-specific pathways by highlighting search's significance in driving traffic.

5.1 Traffic Overview

As noted in Section 1, on average roughly 1 in 40,000 internet users clicked on IRA content on any given day, according to the browser-log dataset, and only 1 in 80,000 clicked on more than one piece of content in a day. We found very small overlap between users exposed to IRA content via both Twitter and Facebook; this is perhaps because the platforms' audiences are different [24], or because there was little cross-promotion between the IRA's Twitter and Facebook campaigns. About 0.02% of the total users in our datasets clicked on content from both campaigns in the same day. The Twitter campaign, we find, received 2.03 times as many clicks from 1.32 as many users as content identified by Facebook.

Table 1 shows the most trafficked Facebook groups and Twitter handles in each campaign, which account for more than 3/4th of total observed traffic in each case. The top-clicked Facebook groups seem to show a mix of topics and political leanings.

Figures 3a and 3b summarize the common referral pathways to IRA properties. Traffic to Facebook-advertised URLs and groups came mainly from Facebook and Search. In fact, we find that traffic from search-engines was a significant source of traffic, accounting for 23.49% of all clicks on Facebook-advertised URLs. The Twitter campaign drew traffic mainly from Twitter, Reddit and Facebook, with Search playing a smaller role. Overall, search-engines accounted for 5.38% of traffic to tweets.

5.2 Traffic by Content-Type

To explore the emotional and political characteristics of content that received traffic on Bing, Twitter and Facebook, we show the results of our crowdsourcing study.

In Figure 4, we characterize the content of each campaign-channel (IRA Facebook ads, IRA tweets and Bing SERP snippet-text[19] of IRA content) as well the clicks this content received from each of three referrers: Bing, Facebook and Twitter. The top row summarizes the IRA content rated along a political axis. The bottom row summarizes IRA content rated along an emotional axis. The X-axis in all the histograms shows the distribution of content (ads or tweets) soft-assigned to each category on the Y-axis.

We see significant differences between and within the different campaign-channels we studied. Between channels, we note that the IRA Twitter content appears to be more right leaning and emotionally neutral than IRA content that appeared on Bing, while the Facebook IRA campaign appeared more left-leaning and high-emotional than content that appeared on Bing (Subfigures 3a and 3c). Right-leaning and left-leaning content drew roughly equal proportions of traffic between the Facebook and Twitter campaigns, while left-leaning and emotionally neutral content drew more impressions on search than they did on other channels (for all pairwise distributions, χ^2-$test$, $p << .01$).

Within campaigns, we observe from Fig. 4a and Fig. 4c that emotionally neutral, apolitical content were the most frequent snippet-types on Bing (Bing, apolitical vs. other, pairwise: z-$test$, $p << .01$, neutral vs. other, pairwise, z-$test$, $p << .01$). There were more left-leaning and apolitical Facebook ads than right-leaning (Facebook content, Left vs. Right: z-$test$, $p << .01$, Apolitical vs. Right: z-$test$, $p << .01$). The IRA Twitter campaign also has large amounts of apolitical tweets, and had more right-leaning content than left, though the difference was not statistically significant (Twitter Content, Left vs. Right: z-$test$, $p = .37$). However, we see that right-leaning and apolitical Twitter content received significantly more clicks (Fig. 4b) than left-leaning Twitter content (Twitter clicks, Right vs. Left: z-$test$, $p << .01$, Apolitical vs. Left: z-$test$, $p << .01$.). In terms of emotional valence, in the Facebook campaign we observe significantly more high-emotion content than neutral content (z-$test$, $p << .01$). In the Twitter dataset, we observe the opposite: more high-emotion content than low, but fail to find significance (z-$test$, $p = .17$).

[19] This snippet text is auto-generated when the SERP is populated. Note that the content we selected to label in our crowdsourcing study are not based on a random sample of all content produced by IRA properties. As described in Section 3 we chose to label the most viewed tweets and search snippets. (We made this decision under a limited budget to capture a sample representative of exposure). As such, distributions over tweets and snippets should not be read as overall content distributions, but rather distributions over content that received the most exposure.

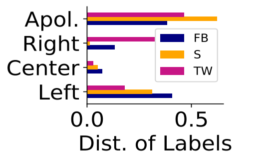

(a) Distribution of political leaning of IRA Facebook posts, tweets and search snippets. We see more left-leaning and apolitical links than right leaning links on Facebook and a large number of apolitical snippets on search. Twitter has roughly equal amounts of right and left content.

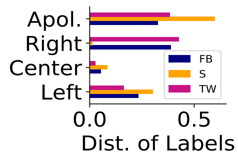

(b) Distribution of clicks across the political-leanings of IRA content by source. Even though there was more left-leaning content on Facebook than right-leaning, the right-leaning content took a larger share of the traffic.

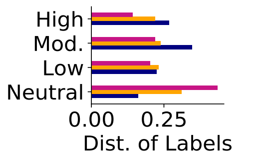

(c) Distribution of content across emotional intensity tags. A large fraction of the IRA promoted content was neural and low emotional intensity. Especially pronounced is the Twitter channel, where nearly 43% of content is rated "Neutral".

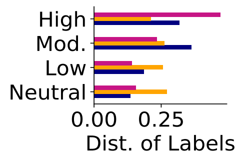

(d) Distribution of clicks across the emotional valence of content by source. Relative to the proportions of overall content produced (c), content with moderate and high emotional valence received more traffic.

Figure 4: Content and Traffic. This figure shows the distribution of the IRA content and traffic across political leaning, emotional intensity and specific politically charged topics using labels tagged in a crowdsourcing experiment. We rated IRA Facebook ads, tweets and search snippets generated for IRA content. The assignment of content to categories was done using the soft-assignment of crowdsourcing labels as described in the methodology section.

6 IRA SEARCH TRAFFIC PATTERNS

In this section we examine attributes of the IRA’s campaign that explain search engine traffic patterns observed in the previous section. We start by describing aspects of the landing pages identified in the Facebook ad dataset. After, we show aspects of the IRA’s Twitter campaign.

6.1 Search-Traffic Characteristics of Facebook-identified Web-Properties

As described in Section 5, apolitical IRA content received a majority of clicks from Bing. Likewise, low and neutral emotion content, and left-leaning content, received a plurality of clicks.

To further characterize Facebook-identified IRA content that performed well on Bing, we study the blackmattersus.com domain. In Figure 5, we split clicks to this domain originating from Bing and those originating from Facebook and find significant differences in the distributions over content-tags (emotions, Facebook vs. Bing: χ^2-$test$, $p = .01$, political-leaning: $fisher$, $p << .01$). As shown, clicks from Facebook are on more emotionally intense and political content (politics, right: z-$test$, $p << .01$, left: z-$test$, $p << .01$, emotions, high: z-$test$, $p << .01$) while clicks from Bing are on more emotionally neutral and apolitical (politics, apolitical: z-$test$, $p << .01$, emotions, neutral: z-$test$, $p << .01$). By publishing neutral, entertaining content the IRA was able to draw additional users to its pages through search engines. We next explore traffic within Bing and find that IRA content was surfaced through preformulated, or auto-populated queries.

6.1.1 *Preformulated Queries.*

While many search-queries in our dataset come from organic queries, i.e. queries where a user directly enters text into the search box, some of the queries also come from auto-populating search boxes with a preformulated query. A user can be lead to an auto-populated search box through many channels including prompts appearing on device lockscreens, and content displayed when opening a new tab in a browser.

We found that IRA content was returned in response to some preformulated queries. Fig. 7 shows that, of all the Facebook-identified content, blackmattersus.com performed especially well through these channels. For example, on March 17$^{\text{th}}$, 2017, as shown in the figure, a link appeared on device lockscreens that took users to the preformulated query: “women scientists." A recently published article entitled “Black Female Scientists. How Many do you Know?”[20] appeared among the search results, thereby capturing impressions.

Of the 6,384 pieces of IRA content published on domains identified via Facebook ads that received impressions on search, 2,725, or 45%, received impressions via preformulated queries. Likewise, preformulated queries accounted for 54% of the overall search impressions of IRA content – but this is primarily due to a few well-performing URLs (indicated by sporadic spikes in Figure 7). When we exclude the top 50 URLs, we find that only 35% of the remaining impressions are mapped to preformulated queries.

As shown in Figure 6, impressions of IRA content via preformulated queries were significantly different than impressions via organic queries (politics, Organic vs. Inorganic, χ^2-$test$, $p << .01$, emotional, χ^2-$test$, $p << .01$). Impressions via preformulated queries were less apolitical and less emotionally neutral than content surfaced in response to organic queries (Apolitical, z-$test$, $p << .01$, neutral, z-$test$, $p << .01$). As shown in Figure 7, content receiving preformulated impressions tied apolitical keywords, like “craigslist” and “lowes”, with politically charged topics. Note that there is not sufficient evidence to conclude that the IRA’s strategy intentionally or systematically exploited preformulated queries as way of information distribution, given the sporadic nature of observed spikes. For example, we fail to observe a measurable increase in the number of impressions via preformulated queries over the time period that we studied. Instead, the few occurrences of such events tend to coincide with timely and popular apolitical search terms.

6.1.2 *Organic Queries.*

46% of impressions received by Facebook-identified IRA content on Bing were received via organic queries. As shown in Table 2, many of these queries were driven by non-news national events: for example, the term “rosa parks”, a relevant topic during black history month, was used in 3% of the organic queries generating impressions of left-leaning IRA URLs. Other queries appear to be news-driven, like queries using the term “14 year old

[20]The blackmattersus.com site is currently no longer operational, but a snapshot of this page can be found at https://web.archive.org/web/20190307234101/https://blackmattersus.com/31081-black-female-computer-scientists-how-many-do-you-know/.

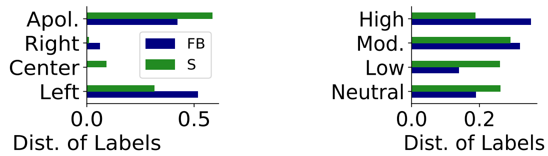

(a) Content-type tags given by crowdsourced workers to blackmattersus.com URLs, weighed by clicks from either Facebook (F) or search (S).

(b) Emotion tags given by crowdsourced workers to blackmattersus .com URLs, weighed by clicks from either Facebook (FB) or search (S).

Figure 5: Case study: Blackmattersus.com, summary. Blac kmattersus.com was the most highly trafficked IRA-owned domain in our dataset. The emotional intensity of pages users reached from Facebook.com is higher than that for pages reached via search, as can be seen by looking at the top-trafficked pages as labeled by crowd workers.

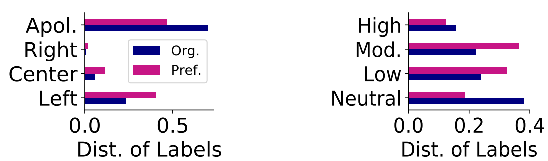

(a) Political leaning of content weighted by impressions from organic (Org) or preformulated (Pref) search channels.

(b) Emotional valence of content weighted by impressions from organic (Org) or preformulated (Pref) search channels.

Figure 6: Organic vs. Preformulated Queries: The IRA URLs that received impressions through preformulated channels tended to be more emotionally intense and less apolitical than content that received impressions in response to organic queries.

latino boy"which generated 11% of the organic impressions of right-leaning IRA URLs. Still other queries have no clear news or time-driven connotation, like those containing the term"patriotic" which was used in queries generating 55% of the organic impressions to right-leaning IRA URLs.[21]

6.1.3 Effects of SERP ordering. IRA content appeared in each of the top five SERP positions in roughly equal proportion for both preformulated and organic queries; we could not observe any statistically significant differences. Clicks were more likely to occur on top results for organic queries: the click-through rate (CTR) for IRA content when it appeared in the first position, for instance, was 4.2 times greater than the CTR for IRA content appearing in the fifth position. This pattern did not hold for preformulated queries. For preformulated queries, the clicks were more evenly distributed.

[21] A negligible percentage of organic queries surfing Facebook-identified IRA properties were navigational ($< .01\%$). We contrast this with our findings for Twitter, and describe our methodology for identifying navigational queries in section 6.2.

6.2 Search-Traffic Characteristics of IRA Tweets

The IRA's Twitter campaign differs from the Facebook-identified URLs in that we observe a lower proportion of impressions received via preformulated queries: only 17% of IRA Twitter impressions are driven through such channels. However, we note two interesting characteristics. First, overall noun-phrases queried by users in searches surfacing IRA tweets show a degree of familiarity with IRA Twitter properties. The top terms, as shown in Table 3 include "pamela moore" and "crystal johnson", which were identified IRA handles. In fact, we observe that 10.3% of Twitter queries appear to be directly navigational, or substantially and unambiguously similar to known IRA Twitter handles.[22] Secondly, we see a pattern of rapid tweeting that we show in a case study, which appears to capture traffic from developing news stories.

6.2.1 Rapid tweeting and search. The IRA tweeted rapidly and in huge volumes across their accounts, tweeting in some cases thousands of times a day, and multiple times within seconds. The high volume of rapid tweets allowed them to be unwittingly incorporated into legitimate news articles at major Western outlets, an observation independently made by [12]. For example, an article published by msn.com[23] references a tweet by @SouthLoneStar, a confirmed IRA account.

In an example, we show that these tweets performed especially well at drawing users searching for information on a breaking news story. The IRA's TEN_GOP[24] Twitter handle, which accounted for 46% of the clicks received by the IRA's Twitter campaign, is a good case study of how an IRA Twitter account can capture significant traffic by being indexed by a search engine during the very early stages of a news cycle. We show in Figure 8 an example of a news-story where rapid tweets from IRA Twitter account allowed them to capture search engine impressions (i.e. be displayed as a search result). Since the tweets were generated before news outlets were able to write stories about the event, they were returned to users by the search engine.

A news event about Dr. Henry Bello occurred on June 30, 2017. This turned out to be a major news story in New York City, receiving thousands of news articles of coverage [21]. However, the bulk of these articles were published and indexed by the search engine on July 1, 2017. Fifteen of the stories published on July 1st were published by the Associated Press, however the remainder came from local domains, with the top domains being: nydailynews.com (12 stories), denik.cz (12 stories), heraldsun.com.au (11 stories) and news-sentinel.com (10 stories). Because the IRA tweeted about the event around noon, they got indexed on June 30, beating many

[22] To identify navigational queries, we calculated *Jaro* string distance between each query and each handle, and selected matches where $jaro(query, handle) > .87$. We selected the threshold .87 after manual verification, and we scanned the list to identify ambiguities between informational and navigational terms: for example, the queries "Tenn GOP" and "Russian allies" matched the handles @TEN_GOP and @russianallies but the user-intent is sufficiently ambiguous, as the queries contained common terms. When the user-intent was ambigious, we considered the query to be informational rather than navigational.

[23] https://www.msn.com/en-gb/news/uknews/london-attack-woman-wearing-hijab-was-distressed-horrified-photographer-says/ar-BByFoiY

[24] TEN_GOP refers to Tennessee GOP. Many IRA accounts mimicked political accounts [35].

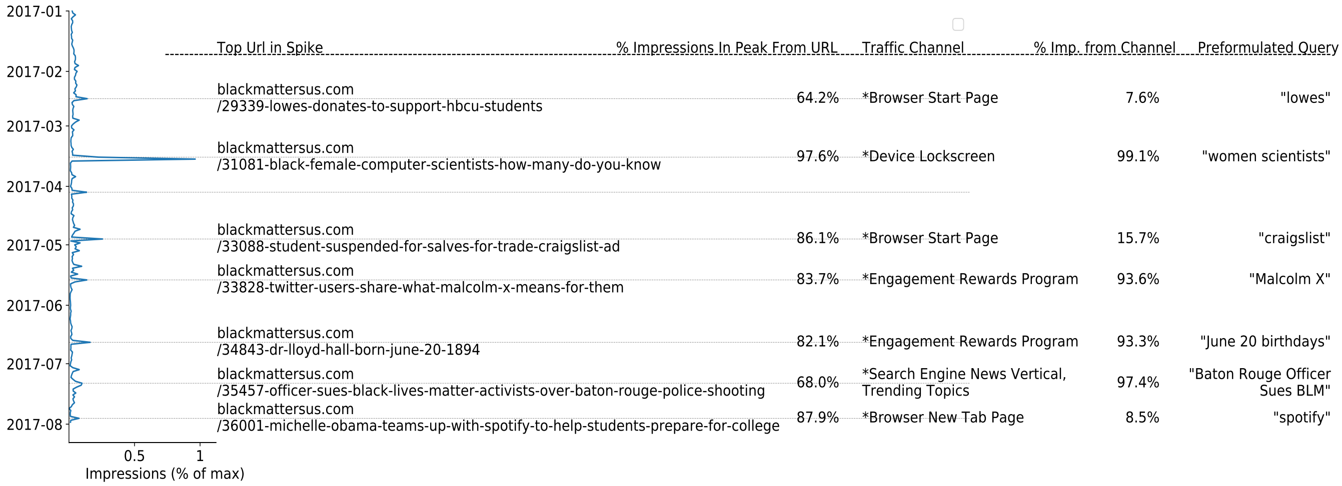

Figure 7: Instances of Facebook-Promoted Domains Appearing in Search Results: The domains most associated with the Facebook-ad campaign often appeared as results to search queries, with surges of search query impressions occurring from time to time. Some of these spikes can be associated with the promotion of various topics. For instance, the largest spike corresponds to the page blackmattersus.com/31081-black-female-computer-scientists-how-many-do-you-know appearing among the results returned for the query, "women scientists" – a query briefly promoted on lock screens of some devices.

Left Leaning	Pct.	Center Leaning	Pct.	Right Leaning	Pct.
rosa parks	0.033	spotify	0.027	patriotic	0.537
trayvon martin project	0.017	kevin hart	0.018	14 year old latino boy	0.119
shea moisture	0.017	tyra banks	0.010	america	0.090

Table 2: Top Search Terms: Facebook. The most common noun-phrases surfacing Facebook-identified IRA URLs via search contain politically-charged terms (e.g. "patriotic") topical terms (e.g. "rosa parks"), and apolitical terms (e.g. "spotify").

Left Leaning	Pct.	Center Leaning	Pct.	Right Leaning	Pct.
resisters	0.054	syria	0.050	tennessee	0.049
katherine johnson	0.029	crystal johnson	0.032	austin hilbourn	0.046
brittney smith twitter	0.020	syria twitter	0.021	pamela moore	0.019

Table 3: Top Search Terms: Twitter. The most common noun-phrases surfacing IRA tweets in queries show a mixture of political topics (e.g. "syria") and known IRA handles (e.g. "pamela moore" and "crystal johnson").

of the articles published and receiving a surge of impressions (i.e. being displayed as a search result) for the story.

7 DISCUSSION

In the current work, we provide a novel set of insights gained by merging fine-grained traffic data with datasets released in the public domain. We build upon active and ongoing research to:

(1) Uncover new IRA tactics that utilized apolitical content to capture traffic on search-engines and receive promotions through indirect promotional channels.
(2) Show how IRA Twitter bots rapidly retweeted local news stories and were able to publish before local news outlets, drawing traffic on search.
(3) Validate observations by other researchers about IRA campaign structure and purpose, and show how internet users interacted with this structure.

We have demonstrated multiple ways by which malicious agents can manipulate the digital landscape, showcasing the large attack surface susceptible to malevolent interventions. We show how search platforms, and their agnostic content-promotion strategies can be leveraged – especially when agents are able to get their late-breaking content indexed before more reputable sources. This is especially the case with Twitter, where bots can respond nearly instantaneously and allow malicious actors to hijack a news cycle.

This work fits into a larger body of ongoing analysis on the IRA campaign conducted by journalists, government officials and academic researchers. Extensive prior work on the IRA has focused on content structure, highlighting connections between different posts cross-platform [12]. A major contribution of this work is the layering of an additional browsing and search datasets onto content structure to validate observations by previous researchers. Many of our observations from these new datasets are novel additions to this conversation: for instance, the IRA's use of apolitical and

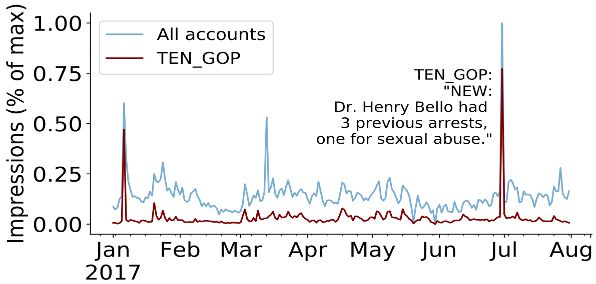

(a) Spike in search impressions (normalized as percent of max) that displayed the TEN_GOP tweet about Henry Bello on June 30, 2017.

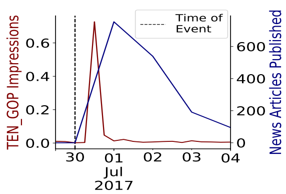

(b) Distribution over impressions received by the TEN_GOP handle, compared with the number of news articles published on the subject.

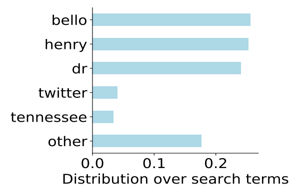

(c) Top search terms leading to impressions received in subfigure (a).

Figure 8: Case study: Tweet impressions for the story of Henry Bello. Subfigure (a) shows impressions on IRA tweets. The surge in impressions is likely linked to the early availability of news information that was unavailable elsewhere; as shown in Subfigure (b) most news outlets only published stories on this subject the following day. Subfigure (c) shows the most common search terms yielding impressions.

neutral content to maximize their reach on search engines and the traffic patterns following this exposure. Additionally, we point to several vulnerabilities, notably the use of emotionally neutral, non-polarizing content that can capture generic search-terms through preformulated queries, and the direct querying of Twitter handles.

Our studies present a set of insights on the strategy, structure, and scope of IRA-related web activities, however, our analyses have limitations, noted in Section 3. While our results are limited by the fact that we study one organization, we do hope that this work can help in predicting and understanding future attacks. Over 70 countries experiencing some form of online disinformation campaign as of 2019 [7]. The insights we gained from merging three major cross-platform datasets suggest an impetus for greater information sharing to identify and respond to emerging threats and tactics.

On a higher level, too, our analysis is confined to and limited by the theoretical framework media critic Todd Gitlin calls an "administrative view on media sociology" [23]. We analyze one set of administrators (the IRA) and the actions taken to extend their reach and engage users, and we analyze another set of administrators (the company we study) and the designs on their properties that allow this to happen. Such measures approximate a vastly more complex system: information diffusion and its effects in society.

To illustrate this limited view, we show in Figure 9 the aggregate number of news articles published about the IRA. As shown, in

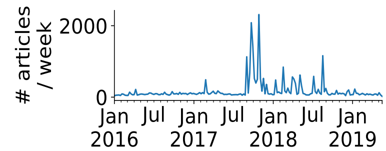

Figure 9: The number of mentions in English-language media articles of the IRA from 2017 to 2019[25] [21]. We observe a large volume of coverage, with peaks of over 2,000 articles published weekly, illustrating the reach of the IRA's campaign extended the traffic directed to the campaign itself.

July 2017, English-language media outlets started reporting on it in heavy volumes, totaling at points 2,000 articles a day. In short, measuring the impact of these campaigns requires further future work. It might be some time until we understand the true ramifications, but an ongoing autopsy of technological vulnerabilities they exploited might, while not suggesting discrete causal actions, help formulate the right direction to investigate.

8 CONCLUSION

We hope this work will further stimulate the development of new approaches to monitoring, understanding, and uncovering propaganda campaigns as well as technology and policy-based solutions to protect important internet infrastructure, and populations engaging with internet services, from political manipulation. Our findings suggest that cross-organization collaboration will be valuable in this endeavor. In future work, we hope to analyze the Senate Select Intelligence Committee's datasets as they become available, herein continuing to provide validation, insights and checks on reported data. We also hope to analyze the second-order reach of the IRA campaign: the user engagement that occurred with news articles and secondary sources commenting on the IRA. Taken together, these directions can give a broader overview of the IRA's campaign and reach and help predict methods for upcoming attacks.

Circumstances around the 2016 U.S. presidential elections and rising concerns about the influence of propaganda campaigns led to the public availability of valuable datasets for understanding IRA activities. Weaving together the public datasets with proprietary data on search and browsing activity provides a previously unavailable lens on the workings of the IRA campaign. We consider this work a preface to numerous opportunities ahead and to the many directions that remain to be explored. We see a dual moving forward, with the Web jointly holding great promise for strengthening liberal democracies while also serving as a platform that can be harnessed by those who seek to manipulate and disrupt. As the digital world continues evolve, and risks to democracy continue to emerge, openness and cooperation among major stakeholders will be essential to understand and counter malevolent threats.

ACKNOWLEDGMENTS

This work was done while Alexander Spangher was an intern and Gireeja Ranade was a researcher at Microsoft Research. The authors would like to thank Michael Golebiewski for sharing insights on the search query analysis and Ryan L. Boyd and James Pennebaker for helpful discussions.

REFERENCES

[1] Aldhous, P. How Russia's Online Trolls Engaged Unsuspecting American Voters. *Buzzfeed* (2018).
[2] Bakshy, E., R. I. C. M. L. A. The Role of Social Networks in Information Diffusion. In *WWW* (2012).
[3] Bessi, A. Personality Traits and Echo Chambers on Facebook. *Computers in Human Behavior* (2016).
[4] Bessi, A., F. E. Social bots distort the 2016 U.S. Presidential election online discussion. *First Monday 21*, 11 (Nov. 2016).
[5] Bittman, L. *The KGB and Soviet disinformation: an insider's view.* Washington: Pergamon-Brassey's, 1985.
[6] Boyd, R., S. A. A. F. B. N. G. R. J. P. E. H. Characterizing the Internet Research Agency's Social Media Operations During the 2016 U.S. Presidential Election using Linguistic Analyses. https://psyarxiv.com/ajh2q, 2018.
[7] Bradshaw, S., and Howard, P. N. *The global disinformation order: 2019 global inventory of organised social media manipulation.* Project on Computational Propaganda, 2019.
[8] Budak, C., A. D. A. A. Limiting the Spread of Misinformation in Social Networks. In *WWW* (2011).
[9] Chen, A. The Agency. *The New York Times* (2015).
[10] Conover, M., R. J. M. F. B. G. F. M. A. F. Political Polarization on Twitter. *ICWSM* (2011).
[11] Conroy, N., R. V. Y. C. Automatic deception detection: Methods for finding fake news. In *ASISST Annual Meeting* (2015), Am. Society for Information Science.
[12] DiResta, R., S. K. B. R. D. S. R. M. R. F. J. A. B. J. The Tactics & Tropes of the Internet Research Agency.
[13] Dutt, R., D. A. E. F. "Senator, We Sell Ads": Analysis of the 2016 Russian Facebook Ads Campaign .
[14] Ebbinghaus, H. *Memory: A Contribution to Experimental Psychology; Translated by Henry A. Ruger and Clara E. Bussenius.* Teachers College, Columbia University, New York, 1913.
[15] Ellick, A. A. W. Disinformation as collaborative work: Surfacing the participatory nature of strategic information operations. *The New York Times.*
[16] Elnoshokaty, A., D. S. D. K. Success factors of online petitions: Evidence from change. org. In *2016 49th Hawaii International Conference on System Sciences (HICSS)* (2016), IEEE, pp. 1979–1985.
[17] Epstein, R., R. R. The search engine manipulation effect (seme) and its possible impact on the outcomes of elections. *Proceedings of the National Academy of Sciences 112*, 33 (2015), E4512–E4521.
[18] Epstein, R., R. R. D. L. C. W. Suppressing the search engine manipulation effect. *Proceedings of the ACM on Human-Computer Interaction 1*, CSCW (2017), 42.
[19] Fourney, A., R. M. G. R. M. M. E. H. Geographic and Temporal Trends in Fake News Consumption During the 2016 US Presidential Election. In *CIKM* (2017), ACM.
[20] Garimella, K., D. F. M. G. A. M. Political discourse on social media: Echo chambers, gatekeepers, and the price of bipartisanship. In *WWW* (2018).
[21] GDELT. Global events database. https://www.gdeltproject.org, 2018.
[22] Gillani, N., Y. A. M. S. V. S. R. D. Me, my echo chamber, and I: introspection on social media polarization. In *WWW* (2018).
[23] Gitlin, T. *Media sociology: the dominant paradigm.* na, 2002.
[24] H., D. A Tale of Two Sites: Twitter vs. Facebook and the Personality Predictors of Social Media Usage, 2012.
[25] Homero, G., J. N. S. V. j. y. p. Social media use for news and individuals' social capital, civic engagement and political participation.
[26] Howard, P., G. B. D. L. J. K. C. F. The IRA, Social Media and Political Polarization in the United States, 2012-2018.
[27] Im, J., C. E. J. S. P. L. T. D. A. B. L. H. D. J. E. G. Still out there: Modeling and identifying russian troll accounts on twitter. *arXiv:1901.11162* (2019).
[28] Inel, O., K. K. C. T. D. A. R. A. P. J. R. L. A. L., and R., S. Crowdtruth: Machine-Human Computation Framework for Harnessing Disagreement in Gathering Annotated Data. In *ISWC* (2014), Springer, pp. 486–504.
[29] Intelligence Community Assessment. Background to "Assessing Russian Activities and Intentions in Recent US Elections": The Analytic Process and Cyber Incident Attribution. Tech. Rep. ICA 2017-01D, Office of the Director of National Intelligence, 01 2017.
[30] Linvill, D., W. P. "The Troll Factories: The Internet Research Agency and State-Sponsored Agenda Building". http://pwarren.people.clemson.edu/Linvill_Warren_TrollFactory.pdf, 2018.
[31] MacFarquhar, N. A powerful russian weapon: The spread of false stories. *The New York Times.*
[32] Mak, T. Russian Influence Campaign Sought To Exploit Americans' Trust In Local News. *National Public Radio* (2018).
[33] Marwick, A., and Lewis, R. Media Manipulation and Disinformation Online. *New York: Data & Society Research Institute* (2017).
[34] Metaxas, T., P. Y. Manipulation of search engine results during the 2016 us congressional elections.
[35] Mueller, R. Internet research agency indictment. https://www.justice.gov/file/1035477/download, 2018.
[36] Mueller, R. Report on the investigation into russian interference in the 2016 presidential election. https://www.justice.gov/storage/report.pdf, 2018.
[37] of State, U. S. D. *Soviet Influence Activities: A Report on Active Measures and Propaganda, 1986-87*, vol. 9627. US Department of State, 1987.
[38] Offer, O. https://www.similarweb.com/, 2018.
[39] Rama, V. Rama, G. K. I. W. A Long-Term Analysis of Polarization on Twitter. In *ICWSM* (2017).
[40] Resnick, P., G. R. T. K. S. M. N. S. Bursting Your (Filter) Bubble: Strategies for Promoting Diverse Exposure. *CSCW* (2013).
[41] Ribeiro, F., S. K. On Microtargeting Socially Divisive Ads: A Case Study of Russia-Linked Ad Campaigns on Facebook, aug 2108.
[42] Schiff, A. Statement on release of facebook advertisements. https://democrats-intelligence.house.gov/news/documentsingle.aspx?DocumentID=379, 2018.
[43] Schiff, A. Statement on release of twitter ads, accounts and data. https://democrats-intelligence.house.gov/news/documentsingle.aspx?DocumentID=396, 2018.
[44] Scott, S. These Are the Ads Russia Bought on Facebook in 2016. *The New York Times* (2017).
[45] Shu, K., S. A. S. W. J. T. H. L. Fake news detection on social media: A data mining perspective. *KDD* (2017).
[46] Starbird, K., A. A. T. W. Disinformation as collaborative work: Surfacing the participatory nature of strategic information operations. *Proceedings of the ACM on Human-Computer Interaction 3*, CSCW (2019), 1–26.
[47] Stewart, L., A. A. K. S. Examining trolls and polarization with a retweet network. In *WSDM Workshop on Misinformation and Misbehavior Mining on the Web.* (2018).
[48] Sunstein, C. *# Republic: Divided Democracy in the Age of Social Media.* Princeton University Press, 2018.