



# Your Friendly Neighborhood Troll: The Internet Research Agency's Use of Local and Fake News in the 2016 US Presidential Campaign

Leon Yin, Franziska Roscher, Richard Bonneau, Jonathan Nagler, and Joshua A. Tucker<sup>1</sup>

## Executive Summary

The NYU Social Media and Political Participation ([SMaPP](#)) lab analyzed data [shared publicly by Twitter](#) on the activity of the Kremlin-linked “[Internet Research Agency](#)” (IRA) to examine whether IRA-operated Twitter accounts spread polarizing or misleading content on social media platforms in an attempt to influence the outcome of the 2016 U.S. presidential election. Our analysis produced 3 main insights:

**Junk News.** IRA-operated accounts shared **50 percent more** junk news than did typical Twitter users, though the overall volume was low. However, the number of junk news sites shared spiked in the weeks immediately leading up to the November 8, 2016 elections.

**Local Media.** The troll accounts shared news from local media sources **15 times more often** than did a comparison group over the same period, possibly trying to take advantage of the [trust](#) local media outlets enjoy in the US. Most of this activity was driven by 27 highly automated IRA accounts that masqueraded as local news sources. However, engagement with these accounts was low, and their activity was not concentrated on content from swing states.

**Polarization.** Several IRA accounts pretended to be American partisans, sharing potentially polarizing local news on Twitter. While engagement with these accounts was higher than with the fake local news outlets, the average number of likes and retweets per link was still **less than 50**.

---

<sup>1</sup> LY wrote the code, conducted the data analysis, and contributed to the first draft of the report. FR contributed to the first draft of the report and subsequent editing. LY, RB, JN & JT contributed to research design, data collection, and editing. The NYU SMaPP Lab is generously funded by the National Science Foundation (SES Award #1756657), the Bill and Melinda Gates Foundation, the John S. and James L. Knight Foundation, the William and Flora Hewlett Foundation, the Rita Allen Foundation, Craig Newmark Philanthropies, and the NYU Global Institute for Advanced Study. Please cite as Yin, Leon, Franziska Roscher, Richard Bonneau, Jonathan Nagler, and Joshua A. Tucker. 2018. “Your Friendly Neighborhood Troll: The Internet Research Agency’s Use of Local and Fake News in the 2016 US Presidential Campaign.” *SMaPP Data Report*. 2018:01.

## Introduction

The “[Internet Research Agency](#)” (IRA), a Kremlin-linked Russian troll farm, is suspected of [having tried to influence the 2016 US presidential elections](#) by spreading polarizing news and [falsehoods online](#). Millions of Americans interacted with IRA accounts on [Twitter](#), [Facebook](#) and other social media platforms in the months leading up to the November 8, 2016 elections without knowing it. The accounts masqueraded as ordinary [Americans](#) or even [local news outlets](#) to spread their messages.

We examined what kinds of links to other websites -- more technically, hyperlinks (URLs) contained in Tweets -- these so-called “troll accounts” shared on Twitter, and, more specifically, what to which types of news content they linked. We also explicitly compared these “link patterns” from the troll accounts to the link patterns of a set of politically interested Twitter users linked, as well as a set of random users of Twitter. We used the domain name of the website where the shared news was published to divide all content sources into three categories: *junk news*, *national mainstream media* and *local mainstream media*.

To do so, our study drew on a dataset shared online by Twitter’s [Elections Integrity](#) Initiative of more than 9 million tweets sent by approximately 3,600 IRA-linked accounts in 2016. Because we are primarily interested in online activity in the US and tweets containing links, we discard tweets without links and accounts that posted in Russian, leaving us with 556 accounts that tweeted approximately 209,000 links between January 2016 and November, 2016.<sup>2</sup>

Data on the comparison groups -- politically engaged users and random users -- were collected by NYU’s Social Media and Political Participation ([SMaPP](#)) lab and cover the same period of activity (January through November 2016). The sample of random Twitter users contains 1,344 accounts that tweeted approximately 106,000 links; the sample of politically engaged users encompasses 1,952 accounts that shared roughly 437,000 URLs.<sup>3</sup>

We found that IRA accounts were more likely than other users to share junk news, but the overall volume of unreliable news sources the trolls linked to was low: six percent of all URLs tweeted by the IRA linked to junk news sites, compared to 4 percent for politically engaged users and 0.4 percent for random users. What stood out, however, was the massive increase in junk links shared in September and October 2018, just prior to the presidential elections.

Perhaps most surprisingly, we found that the IRA relied heavily on local news sources when sharing information online. Thirty percent of all URLs the IRA posted linked to local media outlets. More than half of all IRA accounts under investigation here shared both junk news and local news on Twitter. This could be a concern considering that Americans tend to [trust local media](#) more than any other news

---

<sup>2</sup> See Appendix section Data Description for more details on the dataset. On October 17, 2018, Twitter released the [full collection of Tweets from the IRA accounts](#), which we will be analyzing in the future to investigate the extent to which the conclusions in this report extend to the time before 2016. Thus this report can be considered a preliminary investigation as a means to inform research on the full dataset that is now available.

<sup>3</sup> See Appendix section Data Description for more details on how the data was collected.

source -- the same troll accounts that spread unreliable information online could have also benefited from the credibility of the local news outlets they shared at other times.

We identified 27 IRA accounts that posed as local news outlets on Twitter. These accounts were responsible for 80 percent of all local news content posted by the IRA. Yet engagement with the fake local media outlets was low. Fake local news sources were most active in Kansas, California and Missouri, none of which are swing states.

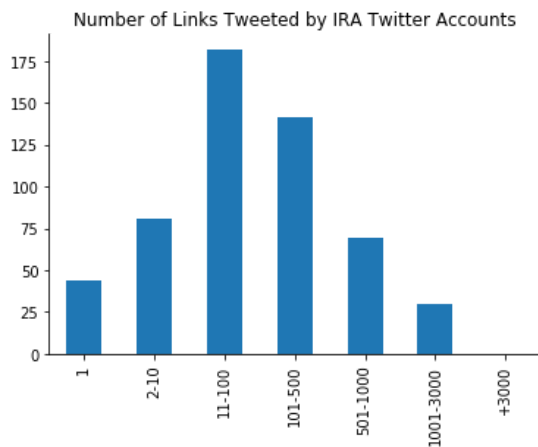
We also identified fake left-leaning and right-leaning partisan accounts. Like the fake local news accounts, the partisan accounts shared local news, but in much fewer links. The local news content posted seemed carefully curated to speak to partisan audiences, talking either about police brutality and the black lives matter movement on the one side, or violence perpetrated by minorities and positive aspects of law enforcement on the other. These partisan accounts also had higher engagement with other Twitter users, with some of their posts amassing hundreds of likes and retweets.

## Account Activity

The 556 IRA-operated accounts under investigation here tweeted a total of around 209,000 links from January 1, 2016 through November 22, 2016. Figure 1 (below) shows that most IRA accounts were far from flooding Twitter with links to online content. Around 8 percent of the 556 troll accounts only shared one URL over the 11 months in question. The majority of accounts posted fewer than 100 URLs in total in 2016.

The top 5 most active accounts are responsible for 23 percent of all hyperlinks shared. All five appear to be fake news outlets, two from Syria, including todayinsyria, an account purporting to be a local Berlin outlet, and two masquerading as U.S. local news outlets, namely DailySanFran and KansasDailynews. Four out of these accounts sent the majority of their Tweets using automated social media managers like [Twitterfeed](#) and [Twibble](#).<sup>4</sup>

**Figure 1: Link Sharing Frequency of IRA Accounts**



<sup>4</sup> For more information on the significance of automation in IRA accounts, see section Local News.

Figure 2 on page 5 traces the weekly URL sharing activity of all IRA-linked accounts combined. The black solid line in the graph shows the total number of links that all IRA accounts together shared per week between January 1, 2016 and November 22, 2016, and the blue dashed-line shows the total number of junk-news sites shared by the IRA-linked accounts over the same period.. The average weekly number of URLs posted during that time is 3,700. As the November 8, 2016 presidential election approached, however, the trolls shared more links: in September 2016, the average weekly number of links jumped to almost 7,000, peaking in October at just under 15,000 links posted in a single week. As we discuss below, these spikes were mostly driven by increases in the number of junk news sites that were shared.

To what kind of content did the IRA link? Of all the URLs shared by IRA accounts, approximately 48 percent led to US news articles of any kind. The other 52 percent linked to content other than US news, like videos on YouTube, social media posts on Facebook or Instagram, or international websites like *sueddeutsche.de*, *dailymail.co.uk* and *skynewsarabia.com*.<sup>5</sup> U.S. news sources can be classified as belonging to one of three categories: *junk news*, *national mainstream media* and *local mainstream media*, as discussed below.

## **Junk News**

One concern voiced about the Internet Research Agency’s online activities was that the troll farm might have spread falsehoods and misleading content online. We checked if and how often IRA-operated accounts shared junk news. The term “junk news” is a value-laden term that means different things to different audiences.<sup>6</sup> To be as precise as possible, we built on the Computational Propaganda Project’s definition of “junk news” to refer to websites that publish false or misleading news.<sup>7</sup> Specifically, we use Merrick College’s OpenSources dataset to identify websites that are known to produce content containing entirely false information, extreme bias, conspiracy theories, hate-based discrimination, clickbait, rumors, state-sponsored news, or junk science.<sup>8</sup>

Junk news stands in contrast to mainstream media, which might occasionally be misinformed while reporting about topics, but generally strive to report accurately, and despite left or right leanings are not considered overly biased.

Perhaps surprisingly, only 6 percent of all URLs shared by IRA accounts led to junk news. This was still considerably more than for the other two groups: just 4 percent of the content posted by politically engaged users led to junk news sites. For the sample of random accounts, the share was a mere 0.4 percent during the same time. That means that even though the overall number of junk content shared was low, IRA-operated accounts tweeted 50 percent more junk news links than the comparison groups in the months leading up to the elections.

---

<sup>5</sup> Links to international websites account for approximately 11.5 percent of the links in the dataset.

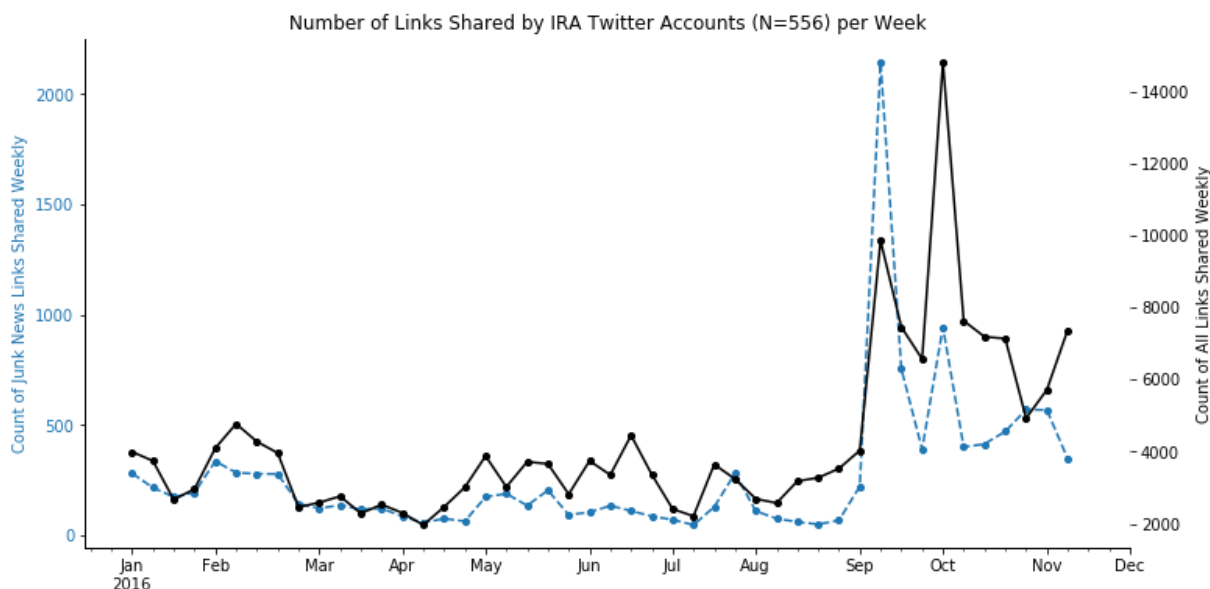
<sup>6</sup> See “Why Do People Share Fake News?” by Alice Marwick, and terminology for problematic information outlined in Caroline Jack’s “Lexicon of Lies” report.

<sup>7</sup> See “Polarization, Partisanship and Junk News Consumption over Social Media in the US” by the Computational Propaganda Group at Oxford.

<sup>8</sup> For examples of junk news sites and their OpenSources tags, see Appendix Table A5.

Figure 2 (below) shows IRA account activity over time: the black line visualizes the average number of links shared by all 556 accounts combined, which hovers around 3,700 per week in the months leading up to the elections, with notable spikes in early September and early October 2016.

**Figure 2:** Overall number of links (black) and number of junk news links (blue) shared by IRA accounts on Twitter.



As the blue line highlights, these jumps in September and October were largely driven by a higher number of links to junk news sites shared during that time. Over the course of 2016, IRA trolls on average shared fewer than 500 junk news links per week. But in September and October 2016, that number soared to up to 2,000 junk news links in a single week, and remained relatively high until after election day.

However, these links to junk news websites were spread by only 51 percent of the troll accounts. That means that **almost half of all IRA-run accounts** only spread reliable information on Twitter. One explanation for this behavior could be that these accounts were so-called “**sleepers**” trying to build a reputation: by posing as legitimate news sources and posting reliable information for months or even years on end, these accounts could potentially accumulate a following and reputation, which could make them more effective tools in any potential future influence campaign.<sup>9</sup> Another explanation is that the trolls could have had goals -- such as exacerbating political polarization -- that could have been accomplished through spreading specifically curated stories from reliable mainstream media sources.

Appendix Table A1 lists the most frequently shared junk news sites for each of the three groups (IRA accounts, politically engaged users and random users). The most frequently shared junk news

<sup>9</sup> In a recent story, [NPR](#) came to a similar conclusion when looking at IRA-operated accounts posing at local news sources on Twitter.

domain among IRA accounts is *Breitbart*, which OpenSources categorizes as biased and unreliable.<sup>10</sup> *Breitbart* is also the most frequently shared junk news domain among politically engaged users, and the third most frequently shared one in the random comparison group.

Other right-leaning junk news sites linked to by the IRA include *The Daily Caller*, a news website co-founded by conservative journalist Tucker Carlson, and Alex Jones' *InfoWars*. *RawStory*, a left-biased news website, also features prominently on the list of links shared by the IRA trolls.

## Mainstream News Sources

In addition to known junk news sites, the IRA accounts also shared links to reliable US media outlets, both national and local. There is reason to suspect that the use of local news sources -- which enjoy relatively high credibility in the US -- might have been strategic, which is why we analyze both categories separately.

### National News Outlets

Overall, **11 percent** of URLs shared by IRA-linked users led to national media outlets. By comparison, the sample of politically interested Twitter users was slightly more likely to link to mainstream news sources: roughly 13 percent of all links shared by this group led to mainstream national news websites. For the random sample, the share was a mere 1.4 percent.

National news sites most frequently shared by IRA-linked accounts include the Congress-focused website *The Hill*, the *Washington Post* and the *Chicago Tribune*. The conservative news outlet *Fox News* (6<sup>th</sup> most popular) was shared more often than the *New York Times* (8th most popular).<sup>11</sup>

The IRA's favorite websites differ from those shared by the comparison groups: Both political users and random users most frequently shared *CNN* content, followed by the *New York Times*. *Fox News* only comes in as the 8th most popular national news outlet for politically engaged users, and ranks 9th for random users.

### Local News Outlets

Perhaps most surprisingly, IRA-operated accounts frequently shared an abundance of local media sources: KSNT station from Topeka, Kansas; San Francisco-based channel KRON4; or regional websites cleveland.com and nj.com are just some of the local outlets frequently linked to by the IRA.<sup>12</sup> Overall, **30 percent** of all links shared by IRA accounts led to local news sources. That means the troll accounts shared **5 times more** local news content than junk content, and 2.5 times more local than national news sources.

The over-time analysis in Figure 3 (below) shows that out of all the news sources the IRA linked to (ignoring entertainment and other content), local news made up more than **60 percent** for most of the

---

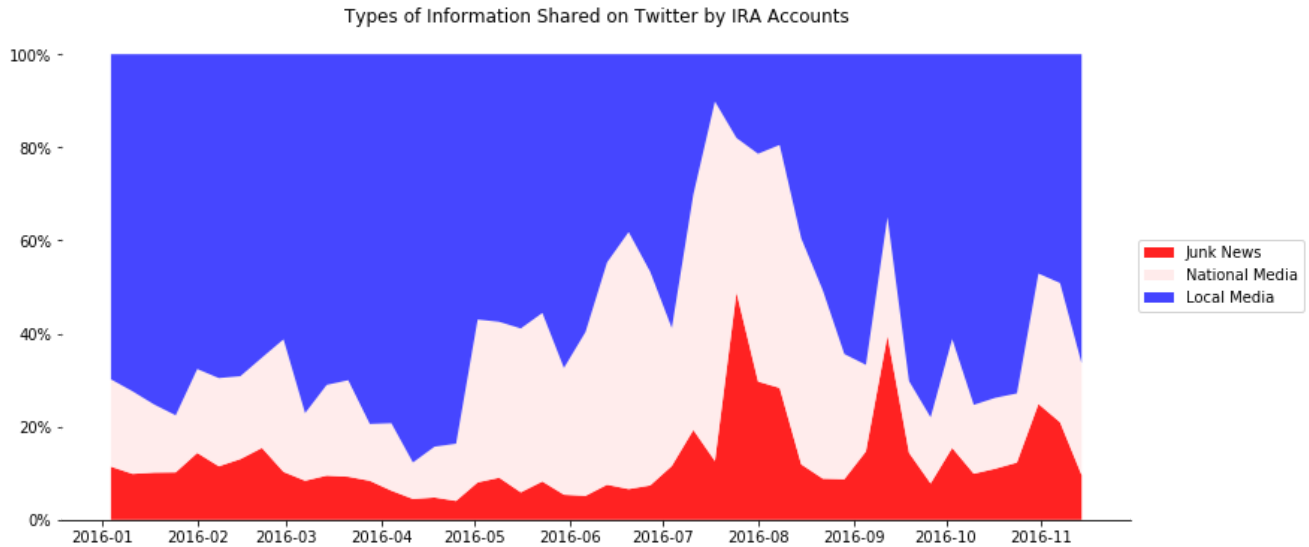
<sup>10</sup> To reiterate, our use of the term "junk news" refers solely to the classification from the OpenSources dataset, and does not reflect any analysis by the SMaPP lab into the quality of the information in these links. For examples of how popular websites were classified by OpenSources, see Appendix Table A5.

<sup>11</sup> See Appendix Table A2 for full list of national news outlets.

<sup>12</sup> See Appendix Table A3 for full list of local news domains.

year. However, in May 2016, the share of national news in the IRA’s news mix rises, from an average of 25 percent to 37 percent from May to September. Starting in August, the proportion of junk news in the news mix briefly spikes. But by election day, local news again make up the majority of the news that the IRA shares.

**Figure 3: Proportion of Each News Category Shared by the IRA over time**



However, not all IRA accounts shared local news sources. Content published by local media outlets was sent by 312 out of 556 accounts (56 percent). Of those 312 trolls, 258 accounts (83 percent) also shared junk news links. That means that **more than half** of all IRA accounts (258 out of 556) shared both local news and junk news over the course of 2016.

One potential reason for the IRA’s local news-sharing activity could be the **enhanced trust** local media outlets enjoy in the US. According to the *Pew Research Center’s 2017 survey on Trust in News Media*, survey respondents trusted local news sources more than national news sources, social media and even more than their friends and family members. A news article by a local media outlet might therefore be perceived as more reliable than an article on the same topic published by a national media outlet. At the same time, sharing trusted sources might help bolster the **credibility** of IRA accounts posing as everyday Americans and local news outlets. Those accounts that shared both junk news and local news might therefore have benefitted from the trust that local news sources enjoy when they posted falsehoods. Similarly, those accounts that only posted local news (and no junk news) might have acted as “**sleeper accounts**,” trying to build a reputation and following, in order to wield more influence online.<sup>13</sup>

<sup>13</sup> For example, Adrian Chen’s [2015 article](#) on the IRA’s #ColumbiaChemicals campaign suggests that the IRA used sleeper accounts pretending to be local news outlets to spread false reports of an ISIS attack on a chemical plant in St. Mary Parish, Louisiana.

Utilizing classifications by researchers at Clemson University,<sup>14</sup> we examined IRA accounts that were posing as local media outlets, as left-leaning US citizens or as right-leaning US citizens.<sup>15</sup> At least 27 accounts (**11 percent** of all accounts) were **masquerading as local media outlets**, calling themselves Atlanta\_Online, TodayMiami, or KansasDailyNews, for example.<sup>16</sup> These 27 fake local news outlets were responsible for the vast majority (80 percent) of all local news sharing by the IRA. Despite the high output, Twitter user engagement with these accounts was rather low: on average, each link sent by these fake local news sources garnered only 0.5 likes. The average retweet rate is not much higher, at 0.8. A content analysis of the tweet accompanying the links shows that these fake local news outlets talked about the presidential election race frequently, with “Trump,” “Clinton” and “#politics” among the most frequently used words. Thus *local news sources* were used by the IRA at least part to spread news about *national politics*.

The fake local media outlets mostly tweeted links to news sources that matched up with their purported **location**. For example, the most active account, KansasDailyNews, shared a total of roughly 17,000 links to local news articles in 2016. Most of these articles came from outlets that are actually located in Kansas (75 percent), or in Kansas City, Missouri (25 percent).

Some IRA accounts also pretended to be **left or right partisans**, based on their Twitter handle and self-descriptions. Fourteen accounts were identified as posing as right-leaning partisans (2.5 percent of all the IRA accounts), using screen names such as USA\_Gunslinger, TheFoundingSon, SouthLoneStar and TEN\_GOP. Based on their self-descriptions, most of these accounts pretended to be white conservatives. They posted 2.3 percent of all local news content the IRA shared. Conversely, accounts posing as left-leaning partisans used screen names such as Cystal1Johnson, TrayneshaCole, BleepThePolice, BlackToLive or blackmattersus; self-descriptions further suggested that the account holders were black Americans.<sup>17</sup> We identified 8 such accounts (1.4 percent of all accounts) that shared only 1 percent of all local news content the IRA posted.

But the fraction of the local news content that these partisan accounts shared was quite polarizing, with a focus on crime, justice, and racial tensions. The right-leaning accounts most frequently tweeted about the “police,” “Trump,” “black,” “guns” and “shootings.”<sup>18</sup> Engagement for these right-leaning accounts was higher than for other IRA accounts: on average, each link they sent out received 19 likes and was retweeted 27 times.<sup>19</sup> Left-leaning accounts shared articles about the “police”, “shootings”, and “#blacklivesmatter”. Engagement with these accounts was even higher than with the right-leaning ones: They got an average of 25 likes per shared link, and were retweeted 47 times on average. The high standard deviations in the analysis suggest that some accounts were extremely successful, while others only elicited minimal engagement. Most surprisingly, more than half of the local news outlets cited by

---

<sup>14</sup> The categorization is taken from Patrick L. Warren and Darren Linvill: “[Troll Factories: The Internet Research Agency and State-Sponsored Agenda Building](#)” (working paper).

<sup>15</sup> We only look at accounts with more than 5,000 followers because when Twitter released the data, they hashed (anonymized) the account IDs for all accounts with fewer than 5,000 followers. Thus we could only apply the classification from Warren and Linvill (2018) to accounts that had more than 5,000 followers.

<sup>16</sup> See Appendix Table A6.

<sup>17</sup> Examples of user self-descriptions in Appendix Table A7.

<sup>18</sup> See Appendix Table A8 for the most frequently used words in texts shared by the three groups (fake local news outlets, fake right-leaning partisans and fake left-leaning partisans).

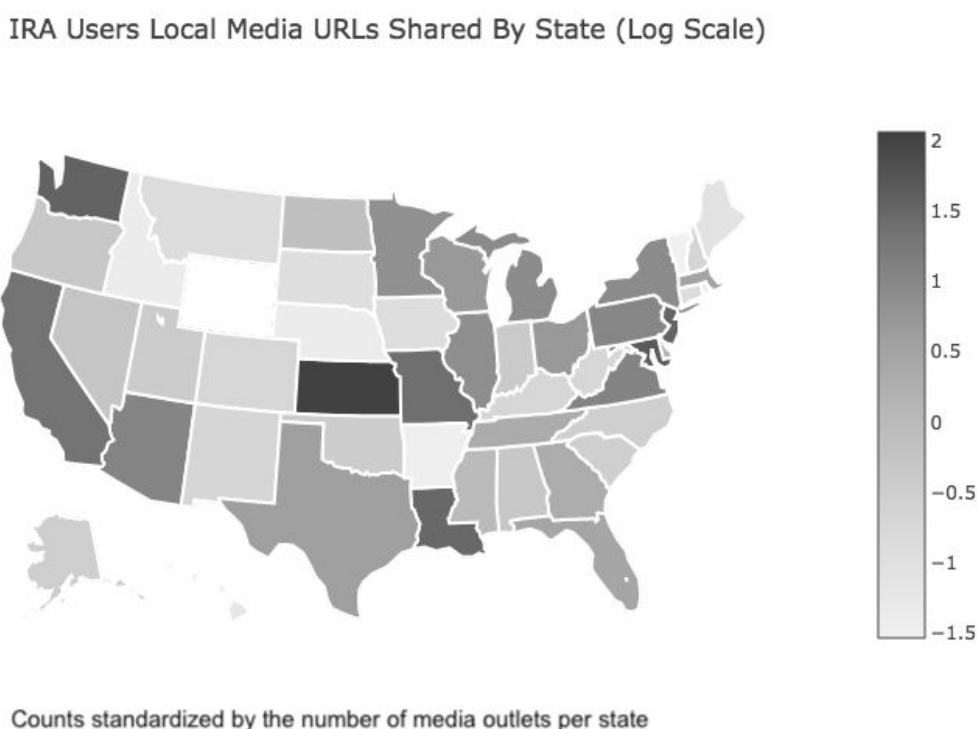
<sup>19</sup> See Appendix Table A9 for engagement metrics.



right-leaning partisans (housed primarily in New York, California, and Texas) are also used by left-leaning partisans. Indeed, how messages are framed and the identity of who is sharing the article enhance and direct polarization in either direction.

Overall, IRA-operated accounts shared the most local news from outlets based in Kansas (17 percent of all local news shared), California (15 percent) and Missouri (7 percent), as Figure 4 shows. Interestingly, none of these were considered swing states in the 2016 elections.<sup>20</sup>

**Figure 4: Map of Local News Shared by IRA Accounts**



Politically engaged users also tweeted local news content, but were less likely to do so: only 1.9 percent of links shared by these users led to local websites, papers or TV stations. When politically engaged users did post local content, a disproportionate amount came from Ohio (12 percent), North Carolina (11 percent), New York State (10 percent) and California (9 percent). While North Carolina and Ohio are traditionally considered swing states, neither New York nor California are. Only 0.33 percent of links in the random sample of users directed to local news sources. Here, most of the local content shared came from Iowa (14 percent, in the past considered a swing state) and Illinois (11 percent, not a swing state).

<sup>20</sup> Missouri, however, had been a swing state in the past, which at least raises the possibility that the IRA might have thought that it was a more politically consequential state for the presidential election than it was.

Overall, 94 percent of content from local media outlets that the IRA shared were posted using social media managers such as Twitterfeed and Twibble.<sup>21</sup> These services provide click statistics, showing how often and when a link was clicked by other users. They also allow for social media activity to be automated to different degrees. Some, such as Hootsuite and Tweetdeck, allow users to compose a post and then schedule it to be published at a later time. Other social media managers, such as Twibble and dlvr.it, for example, make it possible to track websites using an RSS feed, and then tweet content that contains certain keywords. For example, it would be possible to track a news website's RSS feed, and tweet every article that mentions "gun control" or "Donald Trump." Our analysis suggests that the IRA made heavy use of these automation options when posting local content on Twitter.<sup>22</sup>

## Conclusion

An analysis of links shared by Twitter accounts associated with the Kremlin-linked "[Internet Research Agency](#)" (IRA) revealed that the trolls differed in their behavior from politically interested Twitter users and random accounts, the two control groups. Among other things,

- IRA-operated accounts shared **1.5 times** as many links to **junk news** websites, such as *Breitbart*, *Truthfeed* or *Raw Story*, as the most active comparison groups, but the overall share of junk news among all links shared was low (6 percent).
- The percentage of junk news websites the IRA shared **spiked sharply** in September and October 2016, that is, in the weeks immediately leading up to the November 8, 2016 presidential elections.
- The troll accounts heavily banked on **local news** outlets when sharing articles online: 30 percent of all links led to local media content, potentially exploiting the added trust that local news sources enjoy in the US as well as highlighting true events that might be thought to increase political polarization..
- At least 27 IRA accounts **posed as local media** outlets and were responsible for 80 percent of the local news content shared. These accounts heavily relied on social media managers to automate their activity.
- Whether or not a state was a "**swing state**" in the presidential election does not appear to explain whether it was heavily featured in the distribution of local news by IRA accounts.

News articles made up about half of links posted by the IRA. But the troll accounts also shared URLs to other kinds of content: YouTube for example, is within the top 3 domains shared across IRA accounts and comparison groups (Table A4). NYU's SMaPP lab is currently working on analyzing the YouTube content shared by IRA-operated Twitter accounts and will report on this behavior in an upcoming study.

---

<sup>21</sup> Frequency of social media manager use is available in Appendix Table A10 and A11.

<sup>22</sup> For more info on social media managers, see Appendix: *Best practices from Marketing: Link Shorteners and Social Media Managers*

### **Open Source Software and Data**

In conjunction with this report, we are also releasing an open source Python software used to extract and expand URLs -- [urlExpander](#) -- and the [Local News Dataset](#), both of which we used in preparing this report.<sup>23</sup>

### **Acknowledgements**

The SMaPP lab is generously supported by The William and Flora Hewlett Foundation, The Bill and Melinda Gates Foundation, The Rita Allen Foundation, The John S. and James L. Knight Foundation, Craig Newmark Philanthropies, the National Science Foundation, and New York University's Global Institute for Advanced Study and Dean Tom Carew's Research Investment Fund.

---

<sup>23</sup> For more details, see Appendix: *Open Source Tools for Reproducibility*

## Appendix

### Data Description

The Internet Research Agency, or IRA, is a St. Petersburg-based company that operated thousands of social media accounts since 2013. The US Intelligence Community views it as part of Russia's state-run effort to interfere in the 2016 presidential elections, and the firm has been indicted by Special Counsel Robert Mueller.

Twitter so far has identified 3,841 accounts that were maintained by the IRA and tweeted during the 2016 presidential race. More than 1.4 million US-based users interacted with these troll accounts that sent hundreds of thousands of tweets in the weeks leading up to the November 8, 2016 elections.

Twitter released the account IDs to the Senate Intelligence Committee, then suspended them, thereby deleting the tweets from public view. On October 18, 2018, Twitter's Elections Integrity Initiative released a comprehensive archive of these accounts allowing researchers to go beyond case studies to representative results.

We use the Election Integrity IRA archive's 9 millions Tweets which contains IRA accounts that operated in Russian, German, and English environments. For our study, we filter out all accounts that used Cyrillic signs in more than 1% of their tweets, and tweets that link to other tweets. The final IRA dataset used in this report contains 209,156 tweets sent by 556 accounts.

The comparison groups to these IRA-operated accounts are *politically engaged users* and *random users*. Data on these two groups was collected by NYU's SMaPP lab and covers the same period of activity. Politically engaged users are those that tweeted about Donald Trump or Hillary Clinton at least twice during the 2016 presidential election race. The sample contains 1,952 accounts that shared 437,091 URLs. The group of random users is determined by a random number generator that validates if a user with that ID exists on Twitter. The sample contains 1,344 accounts that tweeted 106,416 links. Because the IRA sample and the comparison groups are not of the same size (there are more politically interested and random users than IRA accounts), we only compare *relative* activity, not total activity, across groups.

## **Open Source Tools for Reproducibility**

This study introduces avenues of research by using links as data. To assist the research community to reproduce this research and build upon the methodology, we are open-sourcing the Python software used to extract and expand URLs -- [urlExpander](#). We also believe that the role of local news is underrepresented in the discussions on information operations, media manipulation, and political communication. To facilitate further study of local news ecosystems, we are releasing the [Local News Dataset](#) -- which scrapes several websites into a standardized directory of news organizations. Additional fields in the dataset are: the medium of the news outlet, the state the news outlet is located in, the corporate owners, web domains and social media accounts.

## **Content Classification**

URLs are extracted and expanded from the raw data using the [urlExpander](#) Python package (see Table A4 for the top domains before and after expansion).

To categorize news articles as either *junk news*, *national mainstream news* or *local mainstream news*, we used existing curated lists of media sources. Merrick College's [OpenSources.co](#) was used to classify websites as "junk news." The list bases its classification on mutually non-exclusive categories including overtly false information (fake), state news (state), hate news (hate), rumor mills (rumor), bias, conspiracy theories (conspiracy), clickbait or unreliable content. The [Local News Dataset](#) provides classifications for local mainstream media outlets, and a [list curated by Gregory Eady](#) with the Guardian removed) is the basis for classifying sources as national mainstream media outlets.

## Best Practices from Marketing: Link Shorteners and Social Media Managers

Overall, the majority of accounts analyzed did not post the URLs directly on Twitter, but instead used link shorteners such as bit.ly or ow.ly. Link shorteners transform URLs into short and customizable bitlinks that continue to redirect to the original URL. It is common for publishers to use link shorteners to turn long article URLs into a branded bitlinks (like <https://fxn.ws/2PIrtW2>). Aside from shrinking the link to a more manageable size, link shorteners obfuscate the original source of content and offer analytics.

The top domain shared by the IRA appear to be bit.ly (Table A4). It is only after systematically resolving all URLs, that this study is possible. When a shortened URL is clicked, the link-shortening service records the time, website, device type, and geolocation of the IP address that engaged with the link. Users of link-shortening services can then track the reach of their campaigns on a granular level. In the case of IRA-linked accounts, that means that the IRA could potentially see which links were clicked by their American audience, and where.

In addition to link shorteners, IRA accounts used social media managers like Twibble and Dlv.it.<sup>24</sup> Twibble.io, the fifth most-shared IRA domain, provides a user interface to Twitter's API to automate the delivery of messages that fulfill a user-defined criteria. User-defined criteria fall into one of two categories: *RSS feed* and *Auto Retweet*. *RSS feed* automates posting links directing to news sites, blogs, YouTube channels, Pinterest boards, or any website with an RSS feed. *Auto Retweet* tracks other Twitter accounts and posts messages that contain (or do *not* contain) a set of keywords. These functions operate not only on new content, but offer sampling methods from historical information, such as most popular tweets or random. But social media managers provide more control than simply choosing what kind of content to share. Twibble also allows users to schedule the timing, for example, mimicking work hours in the United States; to set the frequency of posts, such as once every 3 hours; and to limit the number of tweets, for example no more than 100 posts a day. Twibble will also cycle through user-defined "prefixes" or captions, as well as hashtags to make variations of the same post look organic. Twibble provides audience-building strategies by auto-mentioning users who are "top influencers" of hashtags or "top followers" of a given account. Dlv.it provides the same *RSS feed* service and sophisticated control over messaging. However, dlv.it is not limited to Twitter, it also posts on Facebook, Pinterest, Google+, LinkedIn and Tumblr. At the time of this writing, social media managers provide a majority of these functions for free. These services leave a signature either in the text of the message (this is a default setting that can be turned off), or in the traffic source (client name) field of JSON returned by the Twitter API. Using these artifacts, it is possible for everyday users and researchers to screen Twitter accounts for the use of social media managers.<sup>25</sup>

Although Twibble and dlv.it have both these features, the IRA also used social media managers that just offer one of these services like the defunct twitterfeed, which only provided *RSS feed*.

---

<sup>24</sup> [Free video tutorials](#) explaining how social media managers work are available online.

<sup>25</sup> See Table A12 for artifacts left by social media managers.

IRA accounts posing as local media outlets used the *RSS Feed* features prominently. The data shows that sophisticated marketing and branding tools were instrumental to the IRA's activity on Twitter, and potentially to its other social media behavior as well.<sup>26</sup>

---

<sup>26</sup> See the frequency of Tweets and users who used social media managers in Table A10 and A11.



## Tables

**Table A1:** Top Junk news Domains Shared by IRA and Comparison Groups

Top junk news Domains Shared from IRA Twitter Accounts		Top junk news Domains Shared from Political Twitter Accounts		Top junk news Domains Shared from Random Twitter Accounts	
breitbart.com	1826	breitbart.com	1844	lifeneews.com	206
rawstory.com	883	journal-neo.org	1060	wikileaks.org	205
dailycaller.com	686	dailycaller.com	758	libertywritersnews.com	204
americanthinker.com	664	rt.com	741	freebeacon.com	195
truthfeed.com	548	wikileaks.org	714	thefreethoughtproject.com	189
thegatewaypundit.com	494	washingtonexaminer.com	708	wnd.com	171
bb4sp.com	341	dailykos.com	633	therightscoop.com	170
washingtonexaminer.com	299	infowars.com	564	dailykos.com	159
infowars.com	293	thegatewaypundit.com	407	sputniknews.com	150
jihadwatch.org	284	theblaze.com	405	redstate.com	148
rt.com	283	rawstory.com	400	express.co.uk	141
lifeneews.com	206	politicususa.com	358	frontpagemag.com	140
wikileaks.org	205	freebeacon.com	350	zerohedge.com	128
libertywritersnews.com	204	zerohedge.com	336	pjmedia.com	121
freebeacon.com	195	americanthinker.com	326	theblaze.com	121
thefreethoughtproject.com	189	redstate.com	311	therealstrategy.com	119
wnd.com	171	dailywire.com	298	dailysignal.com	108
therightscoop.com	170	conservativetribune.com	242	newsmax.com	93
dailykos.com	159	truthfeed.com	242	cnsnews.com	90
sputniknews.com	150	usuncut.com	198	lifeneews.com	206
redstate.com	148	patriotpost.us	169	wikileaks.org	205
express.co.uk	141	newsmax.com	165	libertywritersnews.com	204
frontpagemag.com	140	therightscoop.com	158	freebeacon.com	195
zerohedge.com	128	commondreams.org	143	thefreethoughtproject.com	189
pjmedia.com	121	twitchy.com	140	wnd.com	171
theblaze.com	121	pjmedia.com	137	therightscoop.com	170
therealstrategy.com	119	endingthefed.com	126	dailykos.com	159
dailysignal.com	108	lifeneews.com	119	sputniknews.com	150
newsmax.com	93	lifezette.com	117	redstate.com	148
cnsnews.com	90	truthrevolt.org	115	express.co.uk	141

**Table A2:** Top National Media Domains Shared by IRA and Comparison Groups

Top National Media Domains Shared by IRA		Top National Media Domains Shared by Political Users		Top National Media Domains Shared by Random Users	
thehill.com	1879	cnn.com	11453	cnn.com	283
washingtonpost.com	1720	nytimes.com	5800	nytimes.com	183
huffingtonpost.com	1213	huffingtonpost.com	3839	huffingtonpost.com	84
reuters.com	1184	washingtonpost.com	3666	washingtonpost.com	74
latimes.com	1157	thehill.com	2049	forbes.com	64
foxnews.com	1059	forbes.com	1845	buzzfeed.com	56
cnn.com	1004	politico.com	1758	wsj.com	48
nytimes.com	996	foxnews.com	1675	vice.com	36
usatoday.com	980	wsj.com	1512	foxnews.com	36
bloomberg.com	853	reuters.com	1509	usatoday.com	33
thedailybeast.com	818	pbs.org	1392	politico.com	33
theroot.com	811	bbc.com	1208	nypost.com	30
chicagotribune.com	637	usatoday.com	984	vox.com	29
politico.com	523	buzzfeed.com	807	businessinsider.com	29
nypost.com	492	politifact.com	708	mic.com	29
townhall.com	376	newsbusters.org	701	bbc.com	27
washingtontimes.com	371	bloomberg.com	613	thehill.com	22
mic.com	366	bostonglobe.com	606	motherjones.com	19
vice.com	359	vox.com	552	slate.com	17
newyorker.com	332	nypost.com	506	theroot.com	17
nbcnews.com	315	vice.com	470	newyorker.com	16
buzzfeed.com	247	latimes.com	426	theatlantic.com	16
motherjones.com	247	theatlantic.com	354	thinkprogress.org	15
wsj.com	241	slate.com	349	qz.com	13
cbsnews.com	200	washingtontimes.com	341	newsweek.com	10
newsone.com	189	nbcnews.com	338	npr.org	10
newrepublic.com	187	motherjones.com	337	nbcnews.com	10
theintercept.com	178	npr.org	299	reuters.com	9
thinkprogress.org	176	nationalreview.com	286	time.com	8

**Table A3:** Top Local Media Domains Shared by IRA and Comparison Groups

Top Local Media Domains Shared from IRA Twitter Accounts		Top Local Media Domains Shared from Political Twitter Accounts		Top Local Media Domains Shared from Random Twitter Accounts	
ksnt.com	9621	whiznews.com	839	desmoinesregister.com	28
kron4.com	5358	nydailynews.com	657	komonews.com	12
seattletimes.com	4981	charlotteobserver.com	458	kcci.com	12
kmbc.com	3282	newsok.com	384	cbslocal.com	11
nj.com	2681	philly.com	199	yakimaherald.com	8
cbslocal.com	2595	cbslocal.com	170	nydailynews.com	8
nola.com	2101	dallasnews.com	166	effinghamdailynews.com	7
ksn.com	2046	newsobserver.com	141	case.edu	7
nydailynews.com	1737	woodtv.com	140	wgntv.com	7
baltimoresun.com	1443	weebly.com	138	wandtv.com	6
startribune.com	1343	mlive.com	103	al.com	6
eastbaytimes.com	1288	wmur.com	88	mlive.com	5
nbccchicago.com	1093	freep.com	81	seattletimes.com	5
kctv5.com	1087	nj.com	78	jg-tc.com	5
wbaltv.com	967	miamiherald.com	72	fox13now.com	4
mysanantonio.com	920	ktla.com	68	sd.gov	4
kwch.com	894	wate.com	59	issuu.com	4
sfgate.com	890	denverpost.com	54	brooklyneagle.com	4
detroitnews.com	873	kansas.com	52	ktla.com	3
nbc12.com	860	dallasobserver.com	52	detroitnews.com	3
azcentral.com	791	wbir.com	50	laweekly.com	3
wafb.com	745	nbcwashington.com	42	postandcourier.com	3
stltoday.com	725	tampabay.com	42	sj-r.com	3
cleveland.com	657	cleveland.com	40	boston.com	3
suntimes.com	596	nola.com	40	fergusfallsjournal.com	3
nbcwashington.com	552	seattletimes.com	39	fox32chicago.com	3
philly.com	539	sfgate.com	38	ktva.com	3
theadvocate.com	519	cbs12.com	36	dl-online.com	3
blogspot.com	470	kcstv.org	34	khou.com	3
post-gazette.com	466	al.com	34	stltoday.com	3

**Table A4:** Top Domains Shared by IRA and Comparison Groups

IRA Users Top Domains		IRA Users Top Domains (Unshortened)	
bit.ly	64968	ksnt.com	9623
sueddeutsche.de	6167	youtube.com	6880
goo.gl	5087	sueddeutsche.de	6195
zeit.de	4472	kron4.com	5358
https://twibble.io	3502	seattletimes.com	4981
fb.me	3283	zeit.de	4525
dlvr.it	3266	facebook.com	3925
youtu.be	3169	twibble.io	3791
ow.ly	2906	kmbc.com	3283
dailym.ai	2664	dailymail.co.uk	3189
youtube.com	2559	sana.sy	2697
vine.co	2201	nj.com	2684
cbsloc.al	2114	cbslocal.com	2595
ln.is	2034	vine.co	2237
instagram.com	1869	instagram.com	2161
buff.ly	1730	nola.com	2106
trib.al	1610	ksn.com	2050
hill.cm	1607	thehill.com	1884
bsun.md	1312	breitbart.com	1830
bayareane.ws	1307	voanews.com	1828
strib.mn	1306	skynewsarabia.com	1781
welt.de	1237	nydailynews.com	1739
payday-loans-24.com	1202	washingtonpost.com	1727
1063.mobi	1198	welt.de	1704
facebook.com	1150	theguardian.com	1618
viid.me	1124	baltimoresun.com	1443
pushpowerpromo.com	1089	startribune.com	1344
ift.tt	1017	blacktolive.org	1311
lat.ms	983	eastbaytimes.com	1290
breitbart.com	977	huffingtonpost.com	1215

SMaPP Data Report 2018:01

Political Users Top Domains		Political Users Top Domains (Unshortened)	
bit.ly	86463	youtube.com	27968
fb.me	32750	facebook.com	15865
youtu.be	18768	instagram.com	14798
instagram.com	13396	cnn.com	11453
dlvr.it	13275	vine.co	9192
cnn.it	10459	fb.me	7054
vine.co	9165	bit.ly	5963
ow.ly	8978	nytimes.com	5800
goo.gl	7854	detik.com	4822
ift.tt	7441	huffingtonpost.com	3839
youtube.com	7069	washingtonpost.com	3666
ln.is	6704	bbc.co.uk	3589
buff.ly	4264	twimg.com	3554
twimg.com	3488	goo.gl	3517
huff.to	3273	thenextweb.com	3216
googling.club	3192	sherif.ws	3215
nyti.ms	3058	stlouisnewschannel.com	3213
sherif.ws	2961	nieuwsblad.be	3201
tinyurl.com	2915	googling.club	3192
tbnn.it	2867	cbc.ca	3173
fbriches.xyz	2706	blubrry.com	2985
trib.al	2169	themexicantimes.mx	2984
j.mp	2117	cubadebate.cu	2882
seekingalpha.com	2069	tbnn.it	2867
sexxbuzz.com	2035	conspirify.com	2755
facebook.com	1982	ntn24.com	2734
wapo.st	1938	fbriches.xyz	2706
sarewah.com	1549	afrika7.com	2348
clubvest.com	1536	people.com	2343
paper.li	1514	sensacine.com	2306

SMaPP Data Report 2018:01

Random Users Top Domains		Random Users Top Domains (Unshortened)	
fb.me	14740	facebook.com	13453
youtu.be	9963	youtube.com	12484
bit.ly	7513	instagram.com	7786
instagram.com	7124	du3a.org	5909
du3a.org	5909	parool.nl	2963
goo.gl	3534	path.com	2818
path.com	2818	vine.co	2362
vine.co	2357	swarmapp.com	2240
swarmapp.com	2240	solomoto.es	2050
youtube.com	1995	is.gd	1879
is.gd	1879	ht.ly	1478
ow.ly	1650	yahoo.co.jp	1434
ht.ly	1478	kuwaitbulletin.com	1375
yahoo.co.jp	1430	twcm.me	1324
kuwaitbulletin.com	1375	tweetedtimes.com	1170
twcm.me	1324	bnent.jp	1028
tweetedtimes.com	1170	jobsindemedia.nl	993
solomoto.es	1096	ameblo.jp	936
bnent.jp	1028	flwrs.com	929
ameblo.jp	931	soundcloud.com	749
flwrs.com	929	twittascope.com	720
buff.ly	800	tumblr.co	696
soundcloud.com	727	sanpedrodelpinatar.es	656
tumblr.co	696	rozee.pk	652
ift.tt	691	twimg.com	649
twimg.com	647	bluesteps.com	638
livedoor.jp	629	liverpoolecho.co.uk	629
liverpoolecho.co.uk	621	livedoor.jp	629
twtslib.com	571	twtslib.com	571
snapd.at	551	jobsindemedia.nl	993

**Table A5:** OpenSources Classifications of Popular Unreliable Websites:

For this study we classify websites that OpenSources.co tagged as bias, clickbait, conspiracy, fake, hate, junksci, rumor, state, and unreliable to be “junk news”. For definitions of each tag, see OpenSource.co. Below are several websites with OpenSources.co’s non-mutually-exclusive classifications.

	<b>bias</b>	<b>clickbait</b>	<b>conspiracy</b>	<b>fake</b>	<b>hate</b>	<b>junksci</b>	<b>rumor</b>	<b>state</b>	<b>unreliable</b>
<b>breitbart.com</b>	1	0	0	0	0	0	0	0	1
<b>dailycaller.com</b>	1	1	0	0	0	0	0	0	0
<b>infowars.com</b>	0	0	1	0	0	0	0	0	0
<b>thegatewaypundit.com</b>	1	0	1	0	0	0	0	0	1
<b>truthfeed.com</b>	1	0	1	0	1	0	0	0	0
<b>rt.com</b>	0	0	0	0	0	0	0	1	0

**Table A6** IRA Accounts Classified as IRA Left, Local, and Right accounts

Screen names of IRA accounts posting as local media outlets and left and right US citizens. Classifications taken from researchers at Clemson University.

IRA Left	IRA Local	IRA Right
BleepThePolice BlackNewsOutlet Crystal1Johnson gloed_up BlackToLive Blk_Voice TrayneshaCole blackmattersus	OnlineCleveland DetroitDailyNew NewOrleansON DailySanFran WashingtOnline DailyLosAngeles TodayBostonMA Seattle_Post SanAntoTopNews ChicagoDailyNew BaltimoreOnline KansasDailyNews PhiladelphiaON OnlineMemphis Atlanta_Online OaklandOnline HoustonTopNews PhoenixDailyNew TodayMiami TodayPittsburgh TodayCincinnati StLouisOnline TodayNYCity ElPasoTopNews DailySanDiego DallasTopNews	SouthLoneStar MissouriNewsUS PigeonToday redlanews USA_Gunslinger patriototus tpartynews TheFoundingSon Pamela_Moore13 Jenn_Abrams rightnpr March_for_Trump TEN_GOP Politweecs

IRA Left user description	IRA Local user description	IRA Right user description
No black person is ugly #BRONZE #BlackLivesMatter #BlackToLive	Breaking news, weather, traffic and more for Atlanta and Georgia. DM us anytime. RTs not endorsements	Business Owner, Proud Father, Conservative, Christian, Patriot, Gun rights, Politically Incorrect. Love my country and my family #2A #GOP #tcot #WakeUpAmerica



**Table A7:** Sample Titles of Articles Shared by IRA Accounts

Below are random samples of Tweets linking to local news that were shared by IRA accounts classified as local news outlets, left-leaning partisans and right-leaning partisans using classifications from Clemson University. The table contains the user screen name, the text (with the link replaced with “\_\_Link\_\_”, and the website that the link directs to.

**IRA Local**

IRA Screen Name	Tweet Text	Link Domain
Seattle_Post	Trace elements of DNA found from Hawaii helicopter crash __LINK__ #news	<a href="http://seattletimes.com">seattletimes.com</a>
KansasDailyNews	Police: Texas teen kills mom, brother before killing himself __LINK__ #news	<a href="http://ksnt.com">ksnt.com</a>
PhiladelphiaON	Philly Fire Station Evacuates After Device Found Outside __LINK__	<a href="http://nbcphiladelphia.com">nbcphiladelphia.com</a>

**IRA Left**

BlackNewsOutlet	Chicago man acquitted of hate crime after slapping, spitting on elderly black judge. __LINK__	<a href="http://suntimes.com">suntimes.com</a>
gloed_up	RT @BleepThePolice: #KillerCops BREAKING!!! 13-Year-Old with replica gun shot by a #Baltimore police officer __LINK__	<a href="http://cbslocal.com">cbslocal.com</a>
Crystal1Johnson	New video shows LAPD officer kick and punch black man lying facedown in the street __LINK__	<a href="http://nydailynews.com">nydailynews.com</a>

**IRA Right**

USA_Gunslinger	Latino activist who said Trump is wrong and illegals are not rapists was charged with rape. __LINK__	<a href="http://fox13now.com">fox13now.com</a>
PigeonToday	Another one for #BlackLivesMatter - look what happens when you don't fight cops __LINK__	<a href="http://wgme.com">wgme.com</a>
redlanews	Islamist w/ arsenal in car taken into custody in Santa Monica! Another attempt of attack! #StopIslam __LINK__	<a href="http://abc7.com">abc7.com</a>

**Table A8:** Top words from Tweets sharing local news aggregated by IRA Left, Local, and Right accounts

IRA Left		IRA local		IRA Right	
police	136	#news	14652	police	13
black	127	man	3510	trump	12
cop	81	police	2998	black	9
man	81	trump	2431	people	9
#btp	53	#politics	2391	gun	8
shot	51	kansas	2252	#tcot	8
officer	46	new	2010	killed	7
#blacklivesmatter	43	#entertainment	1697	shot	7
cops	42	shooting	1365	shooting	7
#caloakland	35	woman	1290	missouri	7
white	32	#business	1229	hillary	7
shooting	31	clinton	1216	cops	7
death	29	killed	1169	lives	6
#policebrutality	28	city	1137	matter	6
school	25	shot	1068	hate	6
killed	25	state	1053	bernie	6
charged	25	says	1043	man	5
woman	24	fire	1011	that's	5
king	22	school	1007	#pjnet	5
jail	21	san	982	don't	5
says	21	crash	961	clinton	5
officers	21	county	914	king	4
arrested	20	#breaking	904	democratic	4
car	19	found	873	victims	4
video	19	arrested	859	boston	4
back	19	video	839	first	4
family	18	dead	825	wrong	4
racist	17	death	813	cop	4
city	17	year-old	803	state	4
men	17	home	799	rape	4

**Table A9:** Engagement metrics for local news articles shared by IRA Left, Local, and Right accounts

For each class of users, the mean and standard deviation for each interaction category (counts for likes, retweets, and replies) are provided. The first number in each cell is the mean (average) number of interactions per account in that group (all, left, local, right), followed by the standard deviation. A large standard deviation indicates that there are large differences between accounts in the same group: some had a lot of interaction while others hardly had any.

	<b>N</b>	<b>likes</b>	<b>retweets</b>	<b>replies</b>
<b>IRA ALL</b>	205,308	1.66 ± 35.16	2.55 ± 28.75	0.19 ± 2.07
<b>IRA Local</b>	49,184	0.53 ± 1.43	0.82 ± 1.68	0.06 ± 0.60
<b>IRA Left</b>	637	25.30 ± 61.22	47.29 ± 116.30	2.04 ± 4.84
<b>IRA Right</b>	144	19.28 ± 44.32	26.84 ± 49.70	3.24 ± 8.28

Although fake local news media are responsible for the bulk (80 percent) of local news sharing, their activity rarely gets engagement. Left- and right-leaning partisans, however, get steady engagement with high standard deviations, suggesting some receive hundreds of retweets. Researchers at University of Washington’s emCOMP Lab also observed the same partisan IRA accounts infiltrating #blacklivesmatter discourse on Twitter.<sup>27</sup>

---

<sup>27</sup> See Arif Ahmer, Leo G. Steward and Kate Starbird’s 2018 paper [“Acting the Part: Examining Information Operations Within #BlackLivesMatter Discourse”](#).

**Table A10:** Client-use Across News Categories

Below are proportions of Tweets sent using each client, as well as proportions of IRA accounts that used each client. Client-use is tied closely to the type of news article shared.

**Client Use of IRA accounts that shared Local Media URLs**

	Tweets	Users
twitterfeed <sup>rss</sup>	77.96%	8.33%
Twibble.io <sup>rss</sup>	7.54%	7.69%
Twitter Web Client	7.51%	90.71%
TweetDeck <sup>sm</sup>	6.97%	23.08%
Twitter for Android	0.02%	0.32%
Buffer <sup>sm</sup>	0.00%	0.32%

Local media articles largely appear to be original content, however we see that the top two sources are both automated social media managers that can automatically siphon content from RSS feeds to Twitter (superscript “RSS”). This helps explain the large quantity of seemingly original local media posts shared from IRA accounts.

**Client Use of IRA accounts that shared National Media**

	Tweets	Users
Twitter Web Client	80.61%	92.84%
TweetDeck <sup>sm</sup>	9.30%	25.34%
twitterfeed <sup>rss</sup>	5.04%	1.93%
Twibble.io <sup>rss</sup>	3.86%	3.31%
dlvr.it <sup>rss</sup>	0.75%	0.28%
Twitter for Android	0.41%	0.28%
Buffer <sup>sm</sup>	0.02%	0.55%

National media is a mix between manual interactions through interfaces like the Twitter Web Client, and automated tools like Twibble, twitterfeed, and dlvr.it. Although most retweets of Junk news are through the Twitter Web Client, a small amount are sourced from RSS feeds using dlvr.it. Some portion of all three types of news are shared using less automated social media dashboards like TweetDeck and Buffer (superscript “SM”). These are used primarily for screening content and scheduling messages.

**Client Use of IRA accounts that shared Junk news**

	Tweets	Users
Twitter Web Client	87.18%	94.39%
TweetDeck <sup>sm</sup>	8.60%	15.44%
dlvr.it <sup>rss</sup>	3.36%	0.35%
Twitter for Android	0.85%	0.35%
Buffer <sup>sm</sup>	0.01%	0.35%

Junk news are mostly from the web client and mostly retweets. This may be an artifact of manually following and monitoring the activity of influential accounts that propagate false and hyper partisan information.

**Table A11** Client Use by IRA, Political, and Random Users

**Clients Used by IRA Accounts**

	Tweets	Users
Twitter Web Client	56.73%	95.50%
twitterfeed <sup>rss</sup>	29.54%	6.47%
TweetDeck <sup>sm</sup>	8.70%	38.31%
Twibble.io <sup>rss</sup>	3.65%	5.22%
dlvr.it <sup>rss</sup>	0.88%	0.18%
Twitter for Android	0.27%	0.18%
Google	0.07%	0.36%
Buffer <sup>sm</sup>	0.05%	0.72%
Medium	0.02%	1.26%
WordPress.com	0.01%	0.18%
Linkis	0.01%	0.18%
Facebook	0.01%	0.18%
IFTTT <sup>rss</sup>	0.01%	0.18%
Twitter Lite	0.01%	0.72%
Commun.it <sup>sm</sup>	0.01%	0.18%

**Clients Used By Random Accounts**

	Tweets	Users
Facebook	15.21%	6.25%
Twitter for iPhone	14.88%	41.71%
Twitter Web Client	12.85%	56.23%
Instagram	8.16%	25.99%
Twitter for Android	7.84%	39.53%
Twitter for Websites	2.84%	32.30%
Hootsuite <sup>sm</sup>	2.76%	2.02%
TweetDeck <sup>sm</sup>	2.43%	3.57%
Google	2.42%	3.88%
twitterfeed <sup>rss</sup>	2.09%	0.57%
Path	1.86%	1.60%
Foursquare	1.70%	3.26%
Tumblr	1.63%	1.65%
iOS	1.53%	17.93%
<a href="http://twittbot.net">twittbot.net</a>	1.48%	0.10%
Twitter for iPad	1.11%	9.72%
تطبيق قرآني	1.10%	0.10%
Twittascope	0.89%	0.21%
Statusbrew <sup>sm</sup>	0.88%	1.50%
fllwrs	0.79%	1.76%
<a href="http://dlvr.it">dlvr.it</a> <sup>rss</sup>	0.78%	0.16%
<a href="http://Ask.fm">Ask.fm</a>	0.76%	3.36%

**Client Use By Political Accounts**

	Tweets	Users
Twitter for iPhone	21.69%	56.77%
Twitter Web Client	16.95%	53.13%
twitterfeed <sup>rss</sup>	13.36%	2.00%
Twitter for Android	10.84%	31.24%
Facebook	5.59%	3.72%
Instagram	3.52%	21.34%
IFTTT <sup>rss</sup>	3.02%	1.18%
Google	2.66%	2.63%
<u>dlvr.it</u> <sup>rss</sup>	2.36%	0.45%
Hootsuite <sup>sm</sup>	2.31%	2.18%
<u>WordPress.com</u>	1.83%	1.45%
Twitter for iPad	1.71%	12.99%
TweetDeck <sup>sm</sup>	1.67%	3.18%
Linkis	1.64%	1.73%
Mobile Web	0.97%	18.98%
Twitter for BlackBerry®	0.93%	0.55%
Buffer <sup>sm</sup>	0.78%	1.00%
Twitter Lite	0.52%	8.45%
RoundTeam <sup>rss</sup>	0.51%	0.18%



**Table A12:** Social Media Managers and the Traces they Leave on Twitter

Social media managers allow users to curate and automate the sharing of content on social media. The following social media managers were discovered from URLs shared by IRA-related Twitter accounts. These services often allow the same user to operate multiple social media accounts on several platforms. The table describes artifacts these social media managers leave behind that are visible to everyday users (“Traces in Text”), and visible programmatically in API responses from the Twitter API (“Traces in the API’s *source* field”). The API response is an HTML a tag. However, in the recent Twitter Elections Integrity dataset, this field is called “tweet\_client\_name” with the HTML stripped away. These are sorted from the most heavily automated information flows at the top, to the most hand-curated information flows at the bottom.

Website	Description	Traces in Text	Traces in the API’s <i>source</i> field	Platforms
Twibble.io	Automated messaging and scheduling of content based on tracked Twitter users and RSS feeds filtered by keywords.	Leaves twibble.io at the end of Tweet text by default. Shortens urls to either twib.in or bit.ly	<a href="http://twibble.io" rel="nofollow">Twibble.io</a>	Twitter, Facebook
dlvr.it	Automated messaging and scheduling of content based on RSS feeds filtered by keywords.	Shortens urls to dlvr.it	<a href="https://dlvrit.com/" rel="nofollow">dlvr.it</a>	Twitter, Facebook
IFTTT	Automated messaging of content based on activity on a tracked social account or event.	Shortens url to ifttt	<a href="https://ifttt.com" rel="nofollow">IFTTT</a>	Twitter, Facebook, YouTube, Spotify, Wordpress, Tumblr, Pinterest
Tweetdeck	Information filtering and tracking. Messaging and scheduling of content based on self-curated content.	None	<a href="http://www.tweetdeck.com" rel="nofollow">Tweet deck</a>	Twitter
Buffer	Information filtering and tracking. Messaging and scheduling of content based on self-curated content.	Shortens url to buff.ly	<a href="https://buffer.com" rel="nofollow">Buffer</a>	Twitter, Facebook, Instagram, Google+, LinkedIn, Pinterest