



Computational
Propaganda
Research Project

Working Paper No. 2017.3

Computational Propaganda in Russia: The Origins of Digital Misinformation

Sergey Sanovich, New York University




Table of Contents

Abstract	3
Introduction.....	3
Domestic Origins of Russian Foreign Digital Propaganda.....	5
Identifying Russian Bots on Twitter	13
Conclusion.....	15
Author Acknowledgements.....	17
About the Author	17
References.....	18
Citation	25
Series Acknowledgement	25

Abstract

Digital propaganda of the Russian government seeks to insulate Putin's leadership from any domestic challengers and aid in his foreign policy ventures, which increasingly set Russian interests off against the West. Yet the propaganda tools, including trolls and bots, were conceived and perfected in the pockets of political competition and a globally integrated market economy still left in Putin's Russia. I discuss how the vibrant Russian blogosphere, left unattended by the government and laser-focused on taking over the traditional media, created the demand for sophisticated online propaganda and censorship tools. I also discuss how the advanced Russian online media and tech sector helped to meet this demand. I conclude with a preliminary report on the detection and exposure of government propaganda online, which could be applicable beyond Russia.

Introduction

The goals and precise impact of the alleged Russian activities around US elections and several other important recent political contests in the West are still subject to a vigorous debate (Kofman, 2017; Bialik & Arthur, 2016; Hopkins, 2016; Enten, 2016; Applebaum, 2016a, 2016b; Applebaum & Lucas, 2016). However, if Russia simply wanted to convince other world powers of its ability to stage a prolonged, multifaceted and global campaign of influence, it has won an unquestionable and impressive victory (Marusic, 2016; Musgrave, 2016). In this campaign, no tool employed by the Russian government attracted more attention than cyber operations of various kinds. Of course, chief among them was the alleged hacking of the email accounts of the Democratic National Committee and Hillary Clinton's campaign chairman. However, all inquiries into the matter emphasize that getting the hacked information across and framing its meaning to the general public was as important as acquiring it in the first place. The United States Intelligence Community report on Russian activities specifically mentions that "Russia's state-run propaganda machine—comprised of its domestic media apparatus, outlets targeting global audiences such as RT¹ and Sputnik,² and a network of quasi-government trolls—contributed to the influence campaign by serving as a platform for Kremlin messaging to Russian and international audiences" (National Intelligence Council, 2017, p. 3). Similarly, a European parliament resolution issued in November 2016 states that "the Russian Government is employing a wide range of

¹ Russian state-owned global TV news network broadcasting in multiple languages, formerly known as Russia Today.

² Russian state-owned international news agency and family of news websites and radio stations, succeeding foreign broadcasting of RIA Novosti and the Voice of Russia, respectively.

tools and instruments, such as think tanks and special foundations (e.g. Russkiy Mir), special authorities (Rossotrudnichestvo), multilingual TV stations (such as RT), pseudo news agencies and multimedia services (e.g. Sputnik), cross-border social and religious groups (...) social media and internet trolls to challenge democratic values, divide Europe, gather domestic support and create the perception of failed states in the EU's eastern neighbourhood" (European Parliament 2016, Section 8).³

The intentions and capability demonstrated by Russia in the domain of cyber propaganda took many Western observers by surprise. Indeed, as opposed to the widely-discussed issue of the builders of the Great Firewall of China (MacKinnon, 2011; King, Pan, & Roberts, 2013), the Russian government prior to 2014 was considered neither particularly artful nor even interested in intervening with the online flows of information (Groves, 2007; see also Kovalev, 2010). In search of an explanation, US defence analysts turned to a 2013 article they discovered in an obscure Russian military–industrial magazine. The article, written by the Chief of the Russian General Staff and General of the Army Valery Gerasimov discussed at length different elements of “ambiguous warfare”, including the information war. Army and navy analysts concluded that online media tools deployed by the Kremlin to brainwash the Ukrainian population and whitewash Russian actions in the West were part of an elaborate strategy clandestinely developed by the Russian military planners (Connell & Evans, 2015; Dickey, Everett, Galvach, Mesko, & Soltis, 2015; cf. McDermott, 2016; Bartles, 2016).

The concern with the alleged “digital propaganda gap”—this New Cold War reincarnation of the original “missile gap”—prompted an intense debate about the impact of and potential responses to the Russian foreign propaganda activities. The European parliament, the British Legatum Institute and the US Center for Strategic and International Studies, among others, published reports on the issue (Conley, Stefanov, Vladimirov, & Mina, 2016; European Union Institute for Security Studies, 2016; Pomerantsev and Lucas, 2016; Russell, 2016). The United States (Timberg, 2016), France (Gramer, 2017) and Germany (Reuters, 2017) have recently appropriated additional funding for various counter-propaganda and cyber-defence measures. Worried about the influence of bots in particular,⁴ the US Defense Advanced Research Projects Agency (DARPA) launched a bot detection programme and ran a competition among scientists to build the best detecting algorithm (Subrahmanian et al., 2016).

It is worth noting that in the broad campaign of the so-called active measures in Eastern

³ This passage is also quoted in the recent House of Commons report on the United Kingdom's relations with Russia, <https://www.publications.parliament.uk/pa/cm201617/cmselect/cmcaff/120/120.pdf>

⁴ The countries of origin of the bots that concerned DARPA are not disclosed in publicly available documents, but the report summarizing the results of the competition mentions Russian bots' activity in Ukraine as a case where the developed methods would be applicable (Subrahmanian et al., 2016, p. 38).

Europe and beyond, Russia is leveraging its traditional advantages: well-trained intelligence and professional diplomatic corps that enjoys unlimited resources, legal immunity and total secrecy at home and can cultivate relationships for decades abroad (Bruno, 2014; Talbott & Brandt, 2017; Remnick, Yaffa, & Osnos, 2017; Snegovaya, 2015). At the same time, the digital elements of the Russian strategy, naturally, had little Soviet foundation to build upon. Neither, as I will show, were they developed by a clandestine effort of Russian military strategists. Instead, their agility and effectiveness were developed through a very long—and very public (rather than secretive)—trial-and-error process. Moreover, if Russia has an edge in digital propaganda, it comes from the most unlikely places in Putin’s Russia: market and political competition. This might have implications for the type of response best suited to effectively counter this type of Russian propaganda.

First, I discuss how the political competition in Putin’s Russia created the demand for online propaganda tools and how market competition in the Russian tech sector was allowed to efficiently meet this demand and create tools that were later deployed in foreign operations. I then discuss some of those tools in detail and describe, with empirical examples, how they could be studied and exposed.

Domestic Origins of Russian Foreign Digital Propaganda

The Atlantic Council’s Digital Forensic Research Lab provides well-documented examples of recent Russian misinformation campaigns, from Armenia to France and from Germany to the United States.⁵ On the one hand, the ability of Russian propaganda to infiltrate dark corners of social media platforms—from the alt-right subreddits to the far-left Twitter threads—with self-serving narratives should not be surprising: this is Russian *modus operandi* in more traditional media too. Expanding on the far-right Cold War tradition of promoting disaffected leftists, RT and Sputnik are skilfully capitalizing on the existing divisions and frustrations in Western societies (Bertrand, 2016; Gorenberg, 2016; Michel, 2017; Michel & Pottier, 2017; Saletan & Carter, 2017; Yablokov, 2015). However, the audience of Russian TV and radio content abroad remains limited (Erickson, 2017) and is probably larger online than offline (Nelson, Orttung, & Livshen 2015), despite significant investments in high-quality production and English-speaking anchors.⁶ In terms of audience, RT is certainly no match for CNN International or BBC World. Information campaigns online, particularly in social media, appear to be much more successful. Why?

⁵ “Three thousand fake tanks” (2017), “Fakes, bots, and blockings in Armenia” (2017), “Hashtag campaign” (2017), “Russian and French Twitter mobs in election push” (2017), “The Kremlin’s audience in France” (2017), “How the alt-right brought #SyriaHoax to America” (2017), “Portrait of a botnet” (2017), “Spread it on Reddit” (2017), “Russian Internet” (2017). See also Broderick (2017).

⁶ “Kremlin boosts budget for TV Channel RT” (2016).

One reason could have its origins in the difference in the domestic media environment that Russian TV broadcasters and social media editors face. The government has virtually monopolized news coverage on all major television channels⁷ and, free from any competition, Russian domestic TV has descended into increasingly crude, evidence-free, often provocative political posturing (Kovalev, 2017). RT, operating in a much more competitive environment, of course, has to adopt more nuanced approaches than its sister channels inside Russia, but this is not what its executives are used to when it comes to broadcasting news on the air.

The Russian social media environment always was—and to a large degree, remains—qualitatively different from the rest of Russian media. Even the people who appear on Russian state-run TV know that they will face much more scrutiny on Twitter and will have to be much more persuasive to get traction there (those who are not prepared to face this level of scrutiny simply avoid social media as they believe that it will be hostile towards them).

This uniquely large difference between media freedom online and offline distinguishes Russia from countries with universally free media and from others, such as China, where both offline and online media were put under tight government control. In Russian domestic politics, this disjointed media environment had important implications for the evolution of Putin's coalition and his relationships with the Russian middle class (Oates & Lokot, 2013). It also prompted the Russian government to adopt bots and trolls as key propaganda tools early on and to gain considerable expertise in deploying them in times of crisis (Chen, 2015).

Vladimir Putin had a chance to appreciate the power of the media to change public opinion and reverse political fortunes at an early stage. The rise to power of a hitherto publicly unknown former KGB lieutenant colonel turned mid-level bureaucrat by the end of the so-called war for Yeltsin's succession was, in large part, the result of a successful media campaign fought on his behalf by a group of Russian oligarchs (Tregubova, 2003, Chapter 10; Gessen, 2012, Chapter 2; Gel'man, Travin, & Marganiya, 2014, pp. 104–108). Putin's media policy in the next 15 years demonstrates that he took this lesson extremely seriously and worked tirelessly to place the media under his control (Burrett, 2010; Lipman, 2009).

Total control of the national TV channels and increasingly tight restrictions put on radio and the print media by the government pushed most serious reporting as well as interested readers to the only area that was left free from interference: the internet. The reasons why Putin treated online media so differently are still unclear. Among the most popular hypotheses are that he hoped to develop an economically profitable

⁷ While all major broadcasters were taken over in the early 2000s, the three smaller and regional independent TV channels, Ren TV, TV Rain (Dozhd) and Tomsk TV2 were put under government control or taken off the air in 2014.

tech sector; that he was concerned about Russia's image abroad, particularly in comparison to that of China (Nossik, 2014); that he wanted to exploit opportunities for online surveillance (Soldatov & Borogan, 2013); and that he viewed it as politically insignificant because of the low internet penetration rate. Indeed, even three years after Putin came to power, in 2002, Russia had 2.1 million people (2 percent of the adult population) who used the internet daily. By 2008 (the end of Putin's second term) this share had increased to 14 million (16 percent of the adult population).⁸ This laissez-faire approach led to a remarkable contrast between traditional and online media. While Freedom House had already downgraded Russia from "Partly Free" to "Not Free" in its annual Freedom of the Press ranking by 2003 (Freedom House, 2003), a monitoring project set up by the Berkman Center for Internet and Society at Harvard noted as late as 2010 that "the Russian blogosphere is a space that appears to be largely free of government control" (Etling et al., 2010, p. 33).

This contrast produced a flourishing online media and tech sector. Their success not only shined against the bleak background of the offline Russian media, but in many respects put it ahead of the curve internationally. In stark contrast with most other countries, Russia's most popular online news media did not represent offline outlets such as newspapers, radio and TV broadcasters. Instead, Gazeta.Ru, Lenta.Ru, NewsRu.com, Polit.ru and the like were built from scratch and became major news producers in their own right. For instance, their staff did original reporting, often as eyewitnesses, instead of simply digitalizing content created by others. Russia is one of the few countries where Google is not the most popular search engine and Facebook is not the most popular social networking website. Remarkably, both occurred without restrictions on American competitors. Unlike Baidu and Weibo, the Russian search engine Yandex and the Russian social media networks Odnoklassniki and Vkontakte won virtually fair competition with their American counterparts.⁹ In a perfect match, a relatively large Russian audience, which quickly regained its economic solvency but maintained highly specific (first and foremost, from the language perspective) content preferences, was well-served by a large pool of well-trained IT professionals, led by a small group of visionary entrepreneurs who decided to seize on the freedoms they got after the collapse of the Soviet Union.

Successful development of local services did not mean that foreign ones were not actively used by Russians. LiveJournal, the most popular Russian social network between 2001 and 2011, despite originally being American and being used predominantly by English speakers, developed a Russian community so large that it was eventually overtaken by a Russian media holding and became dominated by the Russian users

⁸ See <http://bd.fom.ru/report/map/projects/internet/internet1133/vesna2011>

⁹ Still, Vkontakte (but not Yandex or Odnoklassniki) had significantly benefited from the lax enforcement of the property rights. However, this does not make comparison with China less impressive, given that China is also famous for widespread piracy.

(Greenall 2012).

An ample and easily available infrastructure for online communication, in Russian and tailored to local preferences, produced vibrant online news media and a blogosphere that by the end of the 2000s had almost completely supplanted TV and newspapers as the main source of information and platforms for discussing them, at least for educated Russians (Clover 2011). Importantly, the Russian blogosphere set a high bar for the quality of discussions, often featuring original reporting or careful examination of the evidence in the reporting from elsewhere, and it produced many successful opinion leaders (Alexanyan et al., 2012; Etling, Roberts, & Faris, 2014). The impact of the Russian blogosphere was further amplified by the Yandex.Blogs service that featured top blog posts of the day on the main page of the most popular Russian search engine.

While the scale of this activity was still relatively small at the time, the initial decision not to pursue the same strategy of hostile takeover that had already been applied to offline media was a political one and had to do with the power struggle that was taking place within the government rather than with any particular assessment of government as a whole regarding the risks of having a free online press or the challenges in censoring it. Dmitry Medvedev—freshly installed as the third Russian president in May 2008—lacked the power base in the security services that helped Putin so much in entrenching himself in power. He also did not directly control the largest money pools of the government, as those are officially under the jurisdiction of the Council of Ministers, and the new prime minister—Putin himself—was, of course, much more independent than any Russian prime minister before or since. This also diverted large businesses' lobbying efforts from the Kremlin to the prime minister's office. Finally, Medvedev lacked the personal appeal and "street cred" of Putin. In his search for a power base, he identified the emerging Russian middle class—educated professionals, many of them active consumers if not participants in the Russian blogosphere—as the most promising pool of supporters for his re-election (Black, 2014; Black & Johns, 2013; Sakwa, 2014, Chapters 3–5). Largely ignored by the blatant Soviet-style TV propaganda, the middle class appeared to be ripe for more intelligent engagement by a team of enlightened bureaucrats assembled by Medvedev.

Less than a year after assuming office, in early 2009 Medvedev started a video blog which quickly moved to LiveJournal—Russia's main social network and blogging platform at the time. In 2010 he visited Silicon Valley, met Steve Jobs and opened a Twitter account at Twitter headquarters in San Francisco. Notably, his account began to follow (in addition to foreign heads of states and Russian officials) several bloggers known for their criticism of the government and the newsfeed of the radio station Echo of Moscow—perhaps the most critical of government among all the major media outlets in Russia. Finally, in 2011 he started his Facebook page, which he occasionally used to communicate with his readers on matters not covered or covered badly by the

official media (such as the 2011 protests), using a franker tone than in his TV interviews. In all social networks he built up a large readership, which is typical for heads of states, but still notable since the environment was completely different from the general media environment Medvedev was used to: here he could not get his message across simply by eliminating competition and controlling the platform and the agenda (Yagodin, 2012). On a rare occasion in 2011 he visited a small private TV channel, Rain, which at the time was mainly accessible via cable networks and online. As a result, Medvedev became permanently associated with blogging and social networks, and was even called, both in Russia and abroad, “Blogger-in-Chief” (West, 2010). Following Medvedev’s example, several of his aids established a significant presence on social media. In particular, his close aid and economic adviser Arkady Dvorkovich maintains one of the most popular Russian Twitter accounts, with more than 700,000 followers; he also has a Facebook page, as does Medvedev’s press secretary Natalya Timakova (who, as a former journalist with an independent politics and business daily Kommersant, is the Facebook friend of many prominent liberal reporters).

The first known large-scale deployment of pro-government bots and trolls in Russia was carried out in support of this engagement strategy of President Medvedev (Barash & Kelly, 2012; Kelly et al., 2012). The troll contingent was, for the most part, recruited by repurposing pro-Kremlin youth movements, which had been created to combat colour revolution on Moscow’s streets and squares (Hale, 2006). Their job primarily focused on the typical “50-cent-ers”¹⁰ activities: posting diversionary comments in the high-profile opposition blogs (Ananyev & Sobolev, 2017), plus retweeting and reposting pro-government messages.

However, using human trolls for retweeting and reposting is inefficient given that these tasks could easily be automated. Fortunately, by the mid-2000s Russia had a well-established and innovative industry of spam and search optimization.¹¹ Thus originally commercial technology—another child of the flourishing online media and tech industry that developed in Russia without much government interference in the 1990s and 2000s—became a key advantage that the Russian government was able to leverage in its nascent online propaganda strategy.

Meanwhile, trolls, as well as more serious pro-government users, focused on generating content to spread. Following the high bar set by the Russian blogosphere, their posts often featured extensive proofs of their claims. The low-trust environment of Russian

¹⁰ Chinese internet trolls, who are allegedly paid small sums of money for every pro-government post they publish. Recent research suggests that most of them are, instead, permanently on the government payroll, working for various propaganda departments (Miller, 2016). In the Russian case, however, the work was mostly outsourced and the workers were indeed paid per post (Nossik, 2013).

¹¹ One of the early leaders of this industry, German Klimenko, who made a fortune from blog platforms and data analytics, was ultimately appointed Putin’s “Internet adviser” (Turovsky, 2016).

society inhibited reputation building and instead asked any user to prove their point right away, preferably with detailed, often highly technical, reports on the matter. If the point was false or half true, the proof had to be completely or partially faked, but it had to look plausible to have a chance of succeeding. From taming down local property disputes (Balmforth, 2016) to bolstering the incumbent's popularity before a mayoral campaign in Moscow (Suleymanov, 2013), all the way to ensuring the legitimacy of presidential elections (Asmolv, 2014), the ultimate, indisputable proof was needed to win the argument. This rule applied even to pro-government trolls.

Notably, in search of convincing evidence of wrongdoing by the leaders of opposition and independent journalists, the weapon of choice was hacking their emails. A hacker with a fitting nickname, "Torquemada Hell" (later identified as Sergei Maksimov and located and convicted in Germany for his activities), terrorized prominent Russian bloggers with email hacks for years (Tselikov, 2012; Gorbachev, 2015). Information he dug up was then weaponized and spread by bots, trolls and others with the dual goal of compromising victims in the eyes of the general public and sowing discord within the opposition ranks through airing their private, personal grievances against each other in public. Clearly, if the email accounts of the Democratic National Committee or John Podesta were indeed hacked by the Russian government, no additional training was needed to make the best of them.

The trolls' roots in the search optimization industry ensured that early on the focus was not simply on the number of retweets and reposts, but on manipulating search results and popular posts' rankings, targeting engagement not just views. Moreover, even content production was guided by the search optimization algorithms. Analysis by Fedor and Fredheim (2017) of documents leaked by the chief communications manager of pro-Kremlin youth groups reveals her constant obsession with producing content that could climb to the tops of LiveJournal, YouTube and Yandex. The attention paid to the "virality" of the imagery was no less serious than to the political message it carried. Given that LiveJournal and Yandex at the time were virtually free from government control and dominated by users inclined to support the opposition, government propaganda was put to a rigorous test, which has certainly improved its quality, particularly compared to the similar content broadcasted on TV. A similar approach, but one that was carried out in a more systematic fashion, was utilized by RT, when it climbed in the YouTube ratings. As Nelson et al. (2015) show, their channels in all regions and languages contain a significant amount of viral but non-political content (including proverbial cats videos) that draws audience to their political stories.

The opportunity to use existing technologies and independent mechanisms that measure success to achieve government propaganda targets was particularly exciting for the Kremlin officials in charge. Compared to expensive state television, online propaganda was available at a bargain price and allowed the verification of the amount of content produced (by counting, say, the number of retweets) and the

estimation of its impact (by tracking how many hours pro-government posts stayed at the top of LiveJournal or the Russian segment of Twitter). While it has not eliminated misreporting and embezzlement completely (Chen, 2015; Elder, 2012a, 2012b, 2012c), it has probably reduced them in absolute terms (simply because including propaganda in social media is cheaper than using the traditional media) as well as per rouble spent (through feedback and verification mechanisms that are absent offline).

At first, bots were used as a supplementary tool: they were supposed to spread the content produced by trolls and even some genuine pro-government users, who during Medvedev's presidency were (occasionally) willing to engage in discussions with the opposition supporters. When the political environment changed after Putin returned to the Kremlin in 2012—which was accompanied by the largest and longest wave of protests in Russia in two decades (Sakwa, 2014, Chapter 6)—the strategy of engaging with the educated public on social media was deemed a failure along with many other components of Medvedev's presidency (2014, Chapters 7–9). Government propaganda grew cruder and attained clear nationalistic overtones (Laruelle, 2013; Smyth & Soboleva, 2014). In this new environment, bots proved to be a reliable tool, often supplanting trolls and genuine users: when the goal is to block alternative opinions, not engage them in a discussion, easily scalable bot attacks have a natural advantage.

Of course, bots and trolls don't exhaust the menu of options available to the government interested in suppressing alternative views online (see a detailed discussion of the menu in general and Russian government choices in the paper I wrote with Denis Stukal and Joshua Tucker: Sanovich, Stukal, & Tucker 2016). The government could simply filter outlets and platforms it considers threatening. This heavy-handed approach, however, could create significant negative economic consequences, impede the government's own operations and hurt politically neutral and even friendly users, creating a backlash. Alternatively, the government could try to use its legal and market power to influence what kind of media content is created. For example, news websites and blog platforms could be threatened with sizeable fines or even shut down if the user-generated content they host were deemed "extremist" by authorities. This would prompt them to police their content themselves or refrain from hosting it altogether. The government can also prosecute individual bloggers using legal and extra-legal means, as well as take over independent news media and dismiss disloyal editors.

Putin's new government was no longer hesitant about coercing platforms and content producers as well as using technical filtering and distributed denial of service attacks to silence opposition. Persecution of opposition leaders and even ordinary activists increased significantly and became more systematic. New laws were adopted promptly to expand the definition of "extremist views" and to toughen punishment for spreading them (Sakwa, 2014, Chapter 8). However, as with any other authoritarian

government (Howard & Hussain, 2013), it faced the problem of some social media platforms and media outlets being out of reach to both Russian legal regulations and Russian money by virtue of being located outside of Russia. This left Putin with an unenviable choice of shutting them down within Russia completely and bearing all the negative consequences, or letting them remain free so that they provided a powerful platform for alternative opinions. Bots and trolls came in handy in resolving this dilemma and were deployed to deal with web resources that could neither be coerced into policing content on the government's behalf nor bought off. A comparison between domestic Yandex and Vkontakte, on the one hand, and foreign Facebook and Twitter, on the other, illustrates the differential government strategy.

By 2009 the state-owned banking giant Sberbank had already bought the “golden share” of Yandex, the most popular Russian search engine and information portal.¹² The influence that Sberbank had on the company's decision making (coupled with the threat of legal prosecution) paid off when at the height of the Ukrainian crisis Yandex had to close its highly popular ranking service for blogs (Smirnova, 2014).¹³ The most popular Russian social network, Vkontakte (often called the “Russian Facebook”), initially resisted government pressure. When requests to remove pro-Navalny groups came in the wake of large-scale protests after the Duma elections in 2011, Vkontakte owner and CEO, libertarian internet guru Pavel Durov, refused to comply (Razumovskaya, 2011).¹⁴ However, when in early 2014 Vkontakte was served with a request to disclose personal data of the administrators of Euromaidan-related pages on Vkontakte, the government did not take no for an answer. Durov had to sell what was left of his share; he then resigned and left the country (Kononov, 2014; Lunden, 2014, Ries, 2014). Around the same time, right in the middle of the crisis in Crimea, LiveJournal (a Russian company since 2006) had to comply with the government order to permanently ban Alexey Navalny's blog—one of the most popular on the platform.

While domestic providers were relatively easily coerced into enforcing government censorship, global social media platforms proved much more capable of resisting the pressure. For example, in December 2014 the authorities requested that Facebook and Vkontakte block access to pages, allowing supporters of Alexey Navalny to register for a rally protesting against his looming criminal conviction and receive updates about the place and time of the event. Vkontakte blocked the page and all subsequent attempts to create a copy, posting a warning saying, “This page is blocked

¹² The sale, allegedly, took place after negotiations with Dmitry Medvedev and multiple threats to designate companies such as Yandex as “strategic”, which would require them to re-register in Russia and hence severely diminish their appeal on the international capital markets. Yandex is incorporated in the Netherlands as Yandex N.V.—a fact that in 2014 was publicly condemned by Vladimir Putin at his meeting with People's Front for Russia (Brennan, 2014).

¹³ At about the same time, Yandex founder Arkady Volozh had to resign as Yandex's Russian CEO. (He kept the executive position in the international operations, though, see Beard, 2014).

¹⁴ Alexey Navalny is the leading Russian opposition politician. He is subject to multiple ongoing criminal investigations and has spent many months under house arrest. Amnesty International and the Russian human rights group Memorial designated him a prisoner of conscience and a political prisoner, respectively.

upon receiving a Roskomnadzor notification of restricting access to information, which contains calls to participate in mass public events, which fail to follow appropriate regulations, as per the request of the Office of the Prosecutor General of Russia”.¹⁵ Facebook also blocked access to a similar page inside Russia, but after a huge outcry in the Western media, refused to block new copies. Moreover, some Russian media outlets, which were afraid to report the scheduling of the event itself, covered the Roskomnadzor order and the social networks’ response. As a result, more people learned about the event and the new event page that had been started on Facebook attracted even more people.

The Russian authorities had been unable to achieve compliance with selective censorship requests, yet were hesitant to prevent access to platforms like Twitter and Facebook completely. Instead, they deployed bots and trolls to alter the balance of opinions so that they were in their favour (for example, by artificially pushing friendly content and directing users to the readily accessible “trending” lists) and to prevent the use of these platforms for coordination purposes (for example, by flooding the hashtags used by opposition rally organizers with gibberish or counter-propaganda). In the next section I will discuss the preliminary results of the work to identify one particular tool—fully automated bots—on one particular platform, Twitter.

Identifying Russian Bots on Twitter

Focusing on bots in the study of government digital propaganda might not seem very interesting as a scholarly goal: after all, bots by definition do not produce original content, lifting it instead from elsewhere. But this is exactly what has drawn our attention to them: they provide a more direct and clear connection to the owner’s intent. Trolls might tweet based on instructions only some of the time, and provide their own opinions in other tweets. Moreover, drawing the line between a paid troll and a genuine supporter is challenging, and ultimately runs into the question of whether somebody who is not genuinely sympathetic to the government’s cause would do this job. Bots, on the other hand, take all their content from a particular source, and if its content is political, the choice of the source becomes a political decision. It does not mean, of course, that every tweet reflects the political agenda of the owner. On the contrary, as we discovered, many bots post a lot of rather mundane content, such as the entire feeds of major news agencies, which necessarily feature many routine reports that do not have any partisan slant to them. However, the choice of one rather than the other news agency—RT instead of Radio Liberty, for example—is clearly a political choice. Given the amount of data bots can produce, such decisions by a few owners of large bot factories could overshadow individual choices that ordinary users

¹⁵ According to those regulations, authorities could not be notified about the upcoming rally earlier than 15 days in advance. The page was blocked 26 days before the event announced on it was scheduled to take place.

make, thus heavily distorting any impression we might get of what—and who—is popular on Twitter.

In order to distinguish between genuine and automated content generation, together with Denis Stukal and Joshua Tucker we at the Social Media and Political Participation Lab at New York University started collecting Twitter data related to Russian politics in early 2014, when the crisis over Crimea was just starting. The content was automatically downloaded using Twitter API, based on a set of keywords related to Russian politics. The preliminary results I present here are based on more than 14 million tweets posted between February 2014 and December 2015 by more than 1.3 million accounts. Updated results are forthcoming (Stukal, Sanovich, Bonneau, & Tucker, 2016).

Initially, we worried about our ability to find any bots and to come up with a clear definition distinguishing between bots and humans. However, soon after looking into the data we realized that bots are ubiquitous on Russian political Twitter, and are easily identifiable at first glance: they produce a vast number of very similar tweets (for example, post only retweets, or only pictures, or only news headlines) but lack many of the common attributes of human users such as a name, bio, profile picture, location and replies to and from other users, and often (though not always) have no followers. Based on these findings we came up with a simple taxonomy of Twitter accounts (described, with the rest of the methodology, in Sanovich et al. (2016) and in Stukal et al. (2016) and charged a team of human coders to create a training dataset of labelled bots and humans using a very conservative definition of a bot that left any accounts that could possibly belong to a human being or a legitimate organization outside of the bot category. We then used this data set to train a machine learning algorithm to predict the likelihood of any account being a bot, based on a large number of attributes, from the aforementioned number of followers to the frequency of tweeting to the number of hashtags used per tweet (Stukal et al., 2016).

Based on a number of performance metrics, our machine learning algorithm demonstrated very high accuracy in bot identification. Applying it to our data yielded a truly staggering result: among accounts with more than ten tweets in our dataset, around 45 percent are bots.

We also registered a sharp increase in the number of bots (and the amount of content they produce) around the most acute moment in the conflict in Ukraine: in the late summer of 2014, after the downing of the Malaysian Flight 17 and before the new round of fighting (Stukal et al., 2016, Figure 2). This suggests that bots' deployment follows a clear strategy and is well coordinated.

While our analysis of bots' characteristics and behaviour is still underway, one fact is particularly illuminating in the context of the discussion of the evolution of the Russian

government's strategy. While our collection covers an important and consequential moment in recent Russian history, during the conflict with Ukraine and the subsequent period of tumultuous relationships with Western countries, and bots' patterns of activity clearly respond to the conflict dynamics, many of the bot accounts used in this conflict were created years in advance. While we don't have data from that time, it is likely that these accounts were used for purely domestic purposes (for example, against Russian opposition, on behalf of Putin or even Medvedev) before they were deployed to wage a Russian propaganda war in Ukraine and beyond.

Conclusion

Russia could be on a mission to restore its Soviet or imperial glory and to prevent liberal democratic values from taking root in the Russian political system. Yet the tools used are precisely the ones developed in the most internationally competitive part of the Russian economy that emerged during the liberal 1990s and (until recently) was not subject to heavy-handed interventions by the government: the online media and tech sector.

Moreover, tools like bots and trolls were developed for rare cases when the government either wanted to engage opposition in a relatively meaningful discussion online (under Medvedev), or when it wanted to curb it (after Putin came back to the Kremlin in 2012), but was neither able to coerce the foreign platforms hosting the unfriendly discussions to selectively censor them nor prepared to ban these platforms outright.

These external limitations, coupled with the vibrancy and tightness of and the emphasis on the burden of proof in the Russian blogosphere, required the government to build sophisticated tools of online propaganda and counter-propaganda. They combined the ability of bots to jam unfriendly and amplify friendly content and the inconspicuousness of trolls posing as real people and providing elaborate proof of even their most patently false and outlandish claims. The government also utilized existing, independent online tracking and measurement tools to make sure that the content it pays for reaches and engages the target audiences. Last but not least, it invested in the hacking capabilities that allow for the quick production of compromising material against the targets of its smear campaigns.

The latter suggests that building up cyber defence is certainly warranted for electoral campaigns and other entities (including inside Russia) that might become a target of Russian government digital propaganda campaigns. However, the former—the fact that bots and trolls thrive in the low-trust, anything goes, prove-it-on-the-spot environment—also means that building up the reputation of mainstream media, ensuring their objectivity, fairness and professional integrity are trusted by the public,

would do more than anything else to deny Russian digital propaganda the power it currently wields. Beyond that, exposing trolls and bots as well as the nuts and bolts of their campaigns could help both to educate the public in how to avoid falling for the misinformation they spread and to find technological means of disrupting their activity.

Author Acknowledgements

Data presented in Section 3 was collected by the Social Media and Political Participation (SMaPP) Laboratory at New York University (<https://wp.nyu.edu/smapp/>). SMaPP is led by Professors Joshua Tucker, Richard Bonneau, John T. Jost and Jonathan Nagler and is supported by the INSPIRE programme of the National Science Foundation (Award SES-1248077), the New York University Global Institute for Advanced Study, the Moore-Sloan Data Science Environment, and Dean Thomas Carew's Research Investment Fund at New York University.

About the Author

Sergey Sanovich is a Ph.D. candidate in Politics at New York University. He employs experiments, big data and historical sources to study institutions, policies and attitudes that enable authoritarian leaders to come to and stay in power. In the Social Media and Political Participation Lab he investigates censorship and propaganda on social media with a focus on Russia. Sergey holds a B.S. in Economics and an M.S. in Public Policy from the Higher School of Economics (Moscow) and M.A. in Social Sciences from the University of Chicago.

References

- Alexanyan, K., Barash, V., Etling, B., Faris, R., Gasser, U., Kelly, J., Palfrey, J., Roberts, H. (2012). Exploring Russian cyberspace: Digitally-mediated collective action and the networked public sphere. Berkman Center Research Publication No. 2012-2. Retrieved from <http://papers.ssrn.com/abstract=2014998>
- Ananyev, M., & Sobolev, A. (2017, April). Fantastic beasts and whether they matter: Do internet “trolls” influence political conversations in Russia? Paper presented at the meeting of the Midwest Political Science Association. Chicago, IL.
- Applebaum, A. (2016a, December 12). Russia’s next election operation: Germany. The Washington Post. Retrieved from <https://www.washingtonpost.com/news/global-opinions/wp/2016/12/12/russias-next-election-operation-germany/>
- Applebaum, A. (2016b, April 8). The Dutch just showed the world how Russia influences Western European elections. The Washington Post. Retrieved from https://www.washingtonpost.com/opinions/russias-influence-in-western-elections/2016/04/08/b427602a-fcf1-11e5-886f-a037dba38301_story.html?utm_term=.b4dc47e046b8
- Applebaum, A., & Lucas, E. (2016, May 6). The danger of Russian disinformation. The Washington Post. Retrieved from https://www.washingtonpost.com/opinions/the-danger-of-russian-disinformation/2016/05/06/b31d9718-12d5-11e6-8967-7ac733c56f12_story.html
- Asmolov, G. (2014). Kremlin’s cameras and virtual Potemkin villages: ICT and the construction of statehood. In S. Livingstone & G. Walter-Drop (Eds.), *Bits and atoms: Information and communication technology in areas of limited statehood* (pp. 30–46). doi: 10.1093/acprof:oso/9780199941599.001.0001
- Balmforth, T. (2016, February 16). Praise our campaign to destroy the Moscow kiosks and we will reward you. Radio Free Europe/Radio Liberty. Retrieved from <https://www.rferl.org/a/russia-kiosks-destruction-army-of-bots-activists-rewarded/27555171.html>
- Barash, V., & Kelly, J. (2012). Salience vs. commitment: Dynamics of political hashtags in Russian Twitter. Berkman Center Research Publication No. 2012-9. Retrieved from <http://papers.ssrn.com/abstract=2034506>
- Bartles, C. K. (2016). Getting Gerasimov right. *Military Review*, 96(1), 30–38.
- Beard, N. (2014, August 26). Founder and CEO of Yandex, Arkady Volozh, resigns. *Calvert Journal*. Retrieved from <http://calvertjournal.com/news/show/3035/founder-of-yandex-resigns-amid-controversy-arkady-volozh>
- Bertrand, N. (2016, December 10). “A model for civilization”: Putin’s Russia has emerged as “a beacon for nationalists” and the American alt-right. *Business Insider*. Retrieved from <http://www.businessinsider.com/russia-connections-to-the-alt-right-2016-11>
- Bialik, C., & Arthur, R. (2016, November 23). Demographics, not hacking, explain the election results. *FiveThirtyEight*. Retrieved from <https://fivethirtyeight.com/features/demographics-not-hacking-explain-the-election-results/>

- Black, J. L. (2014). *The Russian presidency of Dimitri Medvedev, 2008–2012: The next step forward or merely a time out?* New York: Routledge.
- Black, J. L., & Johns, M. (2013). *Russia after 2012: From Putin to Medvedev to Putin—continuity, change, or revolution?* London: Routledge.
- Brennan, C. (2014, April 24). "Putin Says CIA Created the Internet, Cites Foreign Influence at Yandex." *The Moscow Times*. Retrieved from <http://www.themoscowtimes.com/news/article/putin-says-cia-created-the-internet-cites-foreign-influence-at-yandex/498903.html>
- Broderick, R. (2017, April 25). Here's how far-right trolls are spreading hoaxes about French presidential candidate Emmanuel Macron. *BuzzFeed*. Retrieved from <https://www.buzzfeed.com/ryanhatesthis/heres-how-far-right-trolls-are-spreading-hoaxes-about>
- Bruno, J. (2014, April 16). Russian diplomats are eating America's lunch. *POLITICO Magazine*. Retrieved from <http://www.politico.com/magazine/story/2014/04/russias-diplomats-are-better-than-ours-105773>
- Burrett, T. (2010). *Television and presidential power in Putin's Russia*. London: Routledge.
- Chen, A. (2015, June 2). The agency. *The New York Times*. Retrieved from <http://www.nytimes.com/2015/06/07/magazine/the-agency.html>
- Clover, C. (2011, December 1). Internet subverts Russian TV's message. *The Financial Times*. Retrieved from <https://www.ft.com/content/85dd8e96-1c2d-11e1-9631-00144feabdc0>
- Conley, H., Stefanov, R., Vladimirov, M., & Mina, J. (2016). *The Kremlin playbook: Understanding Russian influence in Central and Eastern Europe*. Washington DC: Center for Strategic/International Studies.
- Connell, M., & Evans, R. (2015). Russia's ambiguous warfare and implications for the U.S. Marine Corps. Retrieved from Center for Naval Analyses website: <http://www.dtic.mil/cgibin/GetTRDoc?Location=U2&doc=GetTRDoc.pdf&AD=ADA618343>
- Dickey, J., Everett, T., Galvach, Z., Mesko, M., & Soltis, A. (2015). Russian political warfare: Origin, evolution, and application. Retrieved from the Naval Postgraduate School website: <https://calhoun.nps.edu/handle/10945/45838>
- Elder, M. (2012a, February 7). Emails give insight into Kremlin youth group's priorities, means and concerns. *The Guardian*. Retrieved from <http://www.theguardian.com/world/2012/feb/07/nashi-emails-insight-kremlin-groups-priorities>
- Elder, M. (2012b, February 7). Hacked emails allege Russian youth group Nashi paying bloggers. *The Guardian*. Retrieved from <http://www.theguardian.com/world/2012/feb/07/hacked-emails-nashi-putin-bloggers>
- Elder, M. (2012c, February 7). Polishing Putin: Hacked emails suggest dirty tricks by Russian youth group. *The Guardian*. Retrieved from <http://www.theguardian.com/world/2012/feb/07/putin-hacked-emails-russian-nashi>
- Enikolopov, R., Petrova, M., & Zhuravskaya, E. (2011). *Media and political persuasion:*

- Evidence from Russia. *American Economic Review*, 101(7), 3253–3285.
- Enten, H. (2016, December 23). How much did Wikileaks hurt Hillary Clinton? *FiveThirtyEight*. Retrieved from <https://fivethirtyeight.com/features/wikileaks-hillary-clinton/>
- Erickson, A. (2017, January 12). If Russia Today is Moscow's propaganda arm, it's not very good at its job. *The Washington Post*. Retrieved from <https://www.washingtonpost.com/news/worldviews/wp/2017/01/12/if-russia-today-is-moscows-propaganda-arm-its-not-very-good-at-its-job/>
- Etling, B., Alexanyan, K., Kelly, J., Faris, R., Palfrey, J. G., & Gasser, U. (2010). Public discourse in the Russian blogosphere: Mapping RuNet politics and mobilization. Berkman Center Research Publication No. 2010-11. Retrieved from <http://papers.ssrn.com/abstract=1698344>
- Etling, B., Roberts, H., & Faris, F. (2014). Blogs as an alternative public sphere: The role of blogs, mainstream media, and TV in Russia's media ecology. Berkman Center Research Publication No. 2014-8. Retrieved from <http://papers.ssrn.com/abstract=2430786>
- European Parliament. (2016). EU strategic communication to counteract anti-EU propaganda by third parties. Resolution, November 23. Retrieved from <http://www.europarl.europa.eu/sides/getDoc.do?pubRef=-//EP//NONSGML+TA+P8-TA-2016-0441+0+DOC+PDF+V0//EN>
- European Union Institute for Security Studies. (2016). EU strategic communications with a view to counteracting propaganda. European Parliament. Retrieved from [http://www.europarl.europa.eu/thinktank/en/document.html?reference=EXPO_IDA\(2016\)578008](http://www.europarl.europa.eu/thinktank/en/document.html?reference=EXPO_IDA(2016)578008)
- Fakes, bots, and blockings in Armenia. (2017). DFRLab. Retrieved from <https://medium.com/dfrlab/fakes-bots-and-blockings-in-armenia-44a4c87ebc46>
- Fedor, J., & Fredheim, R. (2017). "We need more clips about Putin, and lots of them:" Russia's state-commissioned online visual culture. *Nationalities Papers* 45(2), 161–181.
- Freedom House. (2009). *Freedom of the Press 2009: A global survey of media independence*. New York: Rowman & Littlefield Publishers.
- Gel'man, V., Travin, D., & Marganiya, O. (2014). *Reexamining economic and political reforms in Russia, 1985–2000: Generations, ideas, and changes*. Lanham, Maryland: Lexington Books.
- Gessen, M. (2012). *The man without a face: The unlikely rise of Vladimir Putin*. New York: Penguin.
- Gorbachev, A. (2015, July 9). Meet the Hacker Who Terrorized the Russian Blogosphere. *Newsweek*. Retrieved from <http://www.newsweek.com/2015/07/17/gospel-according-hell-351544.html>
- Gorenberg, G. (2016, October 14). The strange sympathy of the far left for Putin. *The American Prospect*. Retrieved from <http://prospect.org/article/strange-sympathy-far-left-putin>
- Gramer, R. (2017, January 10). Wary of Russian cyber threat, France plans to bolster its army of "Digital Soldiers." *Foreign Policy*. Retrieved from <https://foreignpolicy.com/2017/01/10/wary-of-the-russian-cyber-threat-france-plans-to-bolster-its-army-of-digital-soldiers-cyber-attack-europe-elections-hack/>

- Greenall, R. (2012, February 29). LiveJournal: Russia's unlikely internet giant. BBC News. Retrieved from <http://www.bbc.co.uk/news/magazine-17177053>
- Groves, S. (2007). Advancing freedom in Russia. Backgrounder No. 2088. Retrieved from The Heritage Foundation website: <http://www.heritage.org/europe/report/advancing-freedom-russia>
- Hale, H. (2006). Democracy or autocracy on the march? The colored revolutions as normal dynamics of patronal presidentialism. *Communist and Post-Communist Studies*, 39(3), 305–329.
- Hashtag campaign: #MacronLeaks. (2017, May 5). DFRLab. Retrieved from <https://medium.com/dfrlab/hashtag-campaign-macronleaks-4a3fb870c4e8>
- Hopkins, D. (2016, December 20). Voters really did switch to Trump at the last minute. FiveThirtyEight. Retrieved from <https://fivethirtyeight.com/features/voters-really-did-switch-to-trump-at-the-last-minute/>
- How the alt-right brought #SyriaHoax to America. (2017, April 7). DFRLab. Retrieved from <https://medium.com/dfrlab/how-the-alt-right-brought-syriahoax-to-america-47745118d1c9>
- Howard, P., & Hussain, M. (2013). *Democracy's fourth wave? Digital media and the Arab Spring*. Oxford: Oxford University Press.
- Judah, B. (2013). *Fragile empire: How Russia fell in and out of love with Vladimir Putin*. New Haven; London: Yale University Press.
- Kelly, J., Barash, V., Alexanyan, K., Eting B., Faris, R., Gasser, U., & Palfrey, J. G. (2012). Mapping Russian Twitter. Berkman Center Research Publication No. 2012-3. Retrieved from <http://papers.ssrn.com/abstract=2028158>
- King, G., Pan, J., & Roberts, M. (2013). How censorship in China allows government criticism but silences collective expression. *American Political Science Review*, 107(2), 326–343.
- Kofman, M. (2017, January 17). The Moscow school of hard knocks: Key pillars of Russian strategy. *War on the Rocks*. Retrieved from <https://warontherocks.com/2017/01/the-moscow-school-of-hard-knocks-key-pillars-of-russian-strategy/>
- Kononov, N. (2014, March 10). The Kremlin's social media takeover. *The New York Times*. Retrieved from <http://www.nytimes.com/2014/03/11/opinion/the-kremlins-social-media-takeover.html>
- Kovalev, A. (2010, September 24). Russia's blogging revolution. *The Guardian*. Retrieved from <https://www.theguardian.com/commentisfree/2010/sep/24/russia-blogging-revolution>
- Kovalev, A. (2017, March 29). Crippled by the Kremlin, Russia's state media no longer competes. *The Moscow Times*. Retrieved from <https://themoscowtimes.com/articles/crippled-by-the-kremlin-russias-state-media-cant-even-compete-anymore-57577>
- Kremlin boosts budget for TV channel RT. (2016, December 1). *The Moscow Times*. Retrieved from <https://themoscowtimes.com/news/rt-channel-gets-additional-12-bln-rubles-56375>
- Laruelle, M. (2013). Conservatism as the Kremlin's new toolkit: An ideology at the lowest cost. *Russian Analytical Digest*, 138,(8), 2–4.
- Lipman, M. (2009). *Media manipulation and political control in Russia*. Retrieved from Chatham House website:

- <https://www.chathamhouse.org/sites/default/files/public/Research/Russia%20and%20Eurasia/300109lipman.pdf>
- Lunden, I. (2014, April 22). Durov, out for good from VK.com, plans a mobile social network outside Russia. TechCrunch. Retrieved from <http://techcrunch.com/2014/04/22/durov-out-for-good-from-vk-com-plans-a-mobile-social-network-outside-russia/>
- MacKinnon, R. (2011). China's "Networked Authoritarianism." *Journal of Democracy*, 22(2), 32–46.
- Marusic, D. (2016, December 13). What does Russia really want? *The American Interest*. Retrieved from <http://www.the-american-interest.com/2016/12/13/what-does-russia-want/>
- McDermott, R. N. (2016). Does Russia have a Gerasimov doctrine? *Parameters*, 46(1), 97–105.
- Michel, C. (2017, January 13). How Putin played the far left. *The Daily Beast*. Retrieved from <http://www.thedailybeast.com/articles/2017/01/13/how-putin-played-the-far-left.html>
- Michel, C., & Pottier, J.-M. (2017, May 4). The Kremlin's California dream. *Slate*. Retrieved from http://www.slate.com/articles/news_and_politics/foreigners/2017/05/why_russia_cultivates_fringe_groups_on_the_far_right_and_far_left.html
- Miller, B. A. P. (2016). Automated detection of Chinese government astroturfers using network and social metadata. Manuscript in preparation. Retrieved from <https://papers.ssrn.com/abstract=2738325>
- Musgrave, P. (2016, November 28). If you're even asking if Russia hacked the election, Russia got what it wanted. *The Washington Post*. Retrieved from <https://www.washingtonpost.com/posteverything/wp/2016/11/28/whether-or-not-russians-hacked-the-election-they-messed-with-our-democracy/>
- National Intelligence Council. (2017). Assessing Russian activities and intentions in recent US elections. Retrieved from https://www.dni.gov/files/documents/ICA_2017_01.pdf
- Nelson, E., Orttung, R., & Livshen, A. (2015). Measuring RT impact on YouTube. *Russian Analytical Digest*, 177(8), 2–9.
- Nossik, A. (2013, September 10). 11 rubles and 80 kopecks per comment. *Echo of Moscow*. Retrieved from <http://www.echo.msk.ru/blog/nossik/1154616-echo/>
- Nossik, A. (2014, May 15). I helped build Russia's Internet. Now Putin wants to destroy it. *The New Republic*. Retrieved from <http://www.newrepublic.com/article/117771/putins-internet-crackdown-russias-first-blogger-reacts>
- Oates, S., & Lokot, T. (2013). Twilight of the gods?: How the Internet challenged Russian television news frames in the Winter Protests of 2011–12. Manuscript in preparation. Retrieved from <http://papers.ssrn.com/abstract=2286727>
- Pomerantsev, P., & Lucas, E. (2016). *Winning the information war*. Center for European Policy Analysis. London: Legatum Institute.
- Portrait of a botnet. (2017, February 21). DFRLab. Retrieved from <https://medium.com/dfrlab/portrait-of-a-botnet-12fa9d5d6b3>
- Razumovskaya, O. (2011, December 8). Russian social network: FSB asked it to block Kremlin protesters. *The Wall Street Journal*. Retrieved from

- <http://blogs.wsj.com/emerging europe/2011/12/08/russian-social-network-fsb-asked-it-to-block-kremlin-protesters/>
- Remnick, D., Yaffa, J., & Osnos, E. (2017, March 6). Trump, Putin, and the New Cold War. *Annals of Diplomacy, The New Yorker*, 40–55.
- Reuters. (2017, January 9). Germany investigating unprecedented spread of fake news online. *The Guardian*. Retrieved from <https://www.theguardian.com/world/2017/jan/09/germany-investigating-spread-fake-news-online-russia-election>
- Ries, B. (2014, April 16). Founder of “Russia’s Facebook” says government demanded Ukraine protestors’ data. *Mashable*. Retrieved from <http://mashable.com/2014/04/16/vkontakte-founder-fsb-euromaidan/>
- Russell, M. (2016, October). Russia’s information war: Propaganda or counter-propaganda? European Parliamentary Research Service Briefing, European Parliament. Retrieved from [http://www.europarl.europa.eu/thinktank/en/document.html?reference=EPRS_BRI\(2016\)589810](http://www.europarl.europa.eu/thinktank/en/document.html?reference=EPRS_BRI(2016)589810)
- Russian and French Twitter mobs in election push. (2017, April 21). DFRLab. Retrieved from <https://medium.com/dfrlab/russian-and-french-twitter-mobs-in-election-push-bca327aa41a5>
- Russian Internet: Fake news haven? (2017, January 28). DFRLab. Retrieved from <https://medium.com/@DFRLab/russian-internet-fake-news-haven-b5acd9ebd06a>
- Sakwa, R. (2014). *Putin redux: Power and contradiction in contemporary Russia*. London: Routledge.
- Saletan, W., & Carter, P. (2017, March 31). Hate makes us weak. *Slate*. Retrieved from http://www.slate.com/articles/news_and_politics/politica/2017/03/how_russia_capitalizes_on_american_racism_and_xenophobia.html
- Sanovich, S., Stukal, D., & Tucker, J. (2016). Turning the virtual tables: Government strategies for addressing online opposition with an application to Russia. Manuscript submitted for publication. Retrieved from NYU website: https://18798-presscdn-pagely.netdna-ssl.com/smapp/wp-content/uploads/sites/1693/2017/06/Online_Opposition.pdf
- Smirnova, A. (2014, April 18). Yandex.Blogs to partially shut down. *Look At Me*. Retrieved from <http://www.lookatme.ru/mag/live/experience-news/203183-rip-yandex-blogs>
- Smyth, R., & Soboleva, I. (2014). Looking beyond the economy: Pussy Riot and the Kremlin’s voting coalition. *Post-Soviet Affairs*, 30(4), 257–275.
- Snegovaya, M. (2015). Putin’s information warfare in Ukraine: Soviet origins of Russia’s hybrid warfare. Retrieved from <http://understandingwar.org/report/putins-information-warfare-ukraine-soviet-origins-russias-hybrid-warfare>
- Soldatov, A., & Borogan, I. (2013). Russia’s surveillance state. *World Policy Journal*, 30(3), 23–30.
- Spread it on Reddit. (2017, February 10). DFRLab. Retrieved from <https://medium.com/dfrlab/spread-it-on-reddit-3170a463e787>
- Stukal, D., Sanovich, S., Bonneau, R., & Tucker, J. (2016). Detecting Bots on Russian Political Twitter. Manuscript submitted for publication.
- Subrahmanian, V. S., Azaria, A., Durst, S., Kagan, V., Galstyan, A., Lerman, K., . . . Menczer, F.


- (2016). The DARPA Twitter bot challenge. *Computer*, 49(6), 38–46.
- Suleymanov, S. How the Moscow mayoral candidates are winning the Internet. (2013, August 30). *The Interpreter*. Retrieved from <http://www.interpretermag.com/how-the-moscow-mayoral-candidates-are-winning-the-internet/>
- Talbott, S., & Brandt, J. (2017, March 2). What Putin is up to. *The Atlantic*. Retrieved from <https://www.theatlantic.com/international/archive/2017/03/putin-trump-russia-flynn-sessions-hack-kremlin/518412/>
- The Kremlin's audience in France. (2017, April 14). DFRLab. Retrieved from <https://medium.com/dfrlab/the-kremlins-audience-in-france-884a80515f8b>
- Three thousand fake tanks. (2017, January 12). DFRLab. Retrieved from <https://medium.com/@DFRLab/three-thousand-fake-tanks-575410c4f64d>
- Timberg, C. (2016, November 30). Effort to combat foreign propaganda advances in Congress. *The Washington Post*. Retrieved from https://www.washingtonpost.com/business/economy/effort-to-combat-foreign-propaganda-advances-in-congress/2016/11/30/9147e1ac-e221-47be-ab92-9f2f7e69d452_story.html
- Tregubova, Y. (2003). *The tales of a Kremlin digger*. Moscow: Ad Marginem.
- Tselikov, A. (2012, July 23). Russia: Hacker Hell, Scourge of the RuNet. *Global Voices*. Retrieved from <https://globalvoices.org/2012/07/23/russia-hacker-hell-scourge-of-the-runet>
- Turovsky, D. (2016, February 26). Putin's Internet guy: Who is German Klimenko, and how will he advise Russia's president? *Meduza*. Retrieved from <https://meduza.io/en/feature/2016/02/26/putin-s-internet-guy>
- West, D. (2010). President Dmitry Medvedev: Russia's blogger-in-chief. Retrieved from The Brookings Institution website: <http://www.brookings.edu/research/opinions/2010/04/14-medvedev-west>
- Yablokov, I. (2015). Conspiracy theories as a Russian public diplomacy tool: The case of Russia Today (RT). *Politics*, 35(3), 301–315.
- Yagodina, D. (2012). Blog Medvedev: Aiming for public consent. *Europe-Asia Studies*, 64(8), 1415–1434.

Citation

Sergey Sanovich, “Computational Propaganda in Russia: The Origins of Digital Disinformation.” Samuel Woolley and Philip N. Howard, Eds. Working Paper 2017.3. Oxford, UK: Project on Computational Propaganda. comprop.oii.ox.ac.uk<<http://comprop.oii.ox.ac.uk/>>. 32 pp.

Series Acknowledgement

The authors gratefully acknowledge the support of the European Research Council, Computational Propaganda: Investigating the Impact of Algorithms and Bots on Political Discourse in Europe,” Proposal 648311, 2015-2020, Philip N. Howard, Principal Investigator. Additional support has been provided by the Ford Foundation, Google-Jigsaw, and Open Society Foundation. Project activities were approved by the University of Oxford’s Research Ethics Committee. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funders or the University of Oxford.



This work is licensed under a Creative Commons Attribution - Non Commercial - Share Alike 4.0 International License.