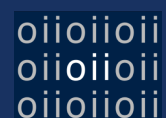# Computational
Propaganda
Research Project

Computational Propaganda Worldwide: Executive Summary

Samuel C. Woolley, *University of Oxford*
Philip N. Howard, *University of Oxford*

UNIVERSITY OF
OXFORD

oiioiioii
oiioiioii
oiioiioii

Table of Contents

Table of Figures

Executive Summary

The Computational Propaganda Research Project at the Oxford Internet Institute, University of Oxford, has researched the use of social media for public opinion manipulation. The team involved 12 researchers across nine countries who, altogether, interviewed 65 experts, analyzed tens of millions posts on seven different social media platforms during scores of elections, political crises, and national security incidents. Each case study analyzes qualitative, quantitative, and computational evidence collected between 2015 and 2017 from Brazil, Canada, China, Germany, Poland, Taiwan, Russia, Ukraine, and the United States.

Computational propaganda is the use of algorithms, automation, and human curation to purposefully distribute misleading information over social media networks. We find several distinct global trends in computational propaganda.

- Social media are significant platforms for political engagement and crucial channels for disseminating news content. Social media platforms are the primary media over which young people develop their political identities.
    - In some countries this is because some companies, such as Facebook, are effectively monopoly platforms for public life.
    - In several democracies the majority of voters use social media to share political news and information, especially during elections.
    - In countries where only small proportions of the public have regular access to social media, such platforms are still fundamental infrastructure for political conversation among the journalists, civil society leaders, and political elites.
- Social media are actively used as a tool for public opinion manipulation, though in diverse ways and on different topics.
    - In authoritarian countries, social media platforms are a primary means of social control. This is especially true during political and security crises.
    - In democracies, social media are actively used for computational propaganda either through broad efforts at opinion manipulation or targeted experiments on particular segments of the public.
- In every country we found civil society groups trying, but struggling, to protect themselves and respond to active misinformation campaigns.

We present new, original evidence about how computational propaganda is produced, managed, and circulated. Each of these cases is important, for different reasons.

- We can measure how Russian Twitter conversation is constrained by highly automated accounts, we can demonstrate how highly automated accounts in the United States moved from peripheral social networks to engagement with core groups of humans, and we trace the source of some forms of junk news and automated accounts to programmers and businesses in Germany, Poland and the United States.
- Interviews with political party operatives, freelance campaigners, and elections officials in seven countries provide evidence that social media bots—and computational propaganda more broadly—have been used to manipulate online discussion.
- Some social media platforms, in particular political contexts, are either fully controlled by or dominated by governments and organized disinformation campaigns. Some 45 percent of Twitter activity in Russia is managed by highly automated accounts. Significant portions of the conversation about politics in Poland over Twitter is produced by a handful of right-wing and nationalist accounts.
- Computational propaganda played a role during three recent political events in Brazil: the 2014 presidential elections, the impeachment of former president Dilma Rousseff and the 2016 municipal elections in Rio de Janeiro.
- The analysis of the social media strategy over Ukraine provides perhaps the most globally advanced case of computational propaganda. Numerous online disinformation campaigns have been waged against Ukrainian citizens on VKontkte, Facebook, Twitter. The industry that drives these efforts at manipulation has been active in this particular country since the early 2000s.
- Authoritarian governments direct computational propaganda at their own population and at populations in other countries. Chinese-directed campaigns have targeted political actors in Taiwan, and Russian-directed campaigns have targeted political actors in Poland and Ukraine.
- In democracies, individual users design and operate fake and highly automated social media accounts. Political candidates, campaigns and lobbyists rent larger networks of accounts for purpose-built campaigns while governments assign public resources to the creation, experimentation and use of such accounts.

- The most powerful forms of computational propaganda involve both algorithmic distribution and human curation—bots and trolls working together. The study of Taiwan reveals that Chinese mainland propaganda over social media is not fully automated but is heavily coordinated.
- We find important examples of positive contributions from algorithms and automation over social media. The Canadian case study reveals some complex algorithms and bots that seek to do constructive public service, though their overall impact is uncertain.

Introduction

A large amount of research has shown that social media plays an important role in the circulation of ideas about public policy and politics. Increasingly, however, social media platforms are conduits for manipulative disinformation campaigns. Political campaigns, governments, and regular citizens around the world are employing both people and bots in attempts to artificially shape public life (Forelle et al., 2015; Woolley, 2016; Gallacher et al., 2017).

Computational propaganda is a term and phenomenon that encompasses recent digital misinformation and manipulation efforts. It is best defined as the use of algorithms, automation, and human curation to purposefully distribute misleading information over social media networks (Woolley & Howard, 2016). Computational propaganda involves learning from and mimicking real people so as to manipulate public opinion across a diverse range of platforms and device networks.

Bots, the automated programs integral to the spread of computational propaganda, are software intended to perform simple, repetitive, robotic tasks. They are used to computationally enhance the ability of humans to get work done online. Social media bots are automated identities that can do mundane tasks like collect information, but they can also communicate with people and systems. They are deployed to do legitimate jobs like delivering news and information. They also are used for more malicious activities associated with spamming and harassment. Whatever their uses, they are able to rapidly deploy messages, interact with other users' content, and effect trending algorithms—all while passing as human users. Political bots, social media bots used for political manipulation, are also effective tools for strengthening online propaganda and hate campaigns. One person, or a small group of people, can use an army of political bots on Twitter to give the illusion of large-scale consensus.

Regimes use political bots, built to look and act like real citizens, in efforts to silence opponents and to push official state messaging. Political campaigns, and their supporters, deploy political bots—and computational propaganda more broadly—during elections in attempts to sway the vote or defame critics. Anonymous political actors harness key elements of computational propaganda such as false news reports, coordinated disinformation campaigns, and troll mobs to attack human

rights defenders, civil society groups, and journalists. Computational propaganda is one of the most powerful new tools against democracy.

Our project at the Oxford Internet Institute at the University of Oxford has built a case-based analysis of computational propaganda in order to better understand its global reach. This is the first systematic collection and analysis of country-specific case studies geared towards exposing and analyzing computational propaganda. We pay particular attention to the themes inherent in propaganda generally but also illuminate crucial details surrounding particular attacks and events. We work to understand who is behind misinformation campaigns while also explaining who the victim groups are, what they experience, and what they—and others fighting this global problem—can do.

Altogether, the authors of these case studies use a broad definition of "social media", using it to refer to (a) the information infrastructure and tools used to produce and distribute content that has individual value but reflects shared values; (b) the content that takes the digital form of personal messages, news, and ideas that becomes cultural products; and (c) the people, organizations, and industries that produce and consume both the tools and the content (Howard, 2011; Howard & Parks, 2012, p. 359).

Computational propaganda flourished during the 2016 US Presidential Election (Howard, Kollanyi, & Woolley, 2016). There were numerous examples of misinformation distributed online with the intention of misleading voters or simply earning a profit. Multiple media reports have investigated how "fake news" may have propelled Donald J. Trump to victory (Dewey, 2016; Parkinson, 2016; Read, 2016). In Michigan, one of the key battleground states, junk news was shared just as widely as professional news in the days leading up to the election (Howard, Bolsover, Kollanyi, Bradshaw, & Neudert, 2017). There is growing evidence that social media platforms support campaigns of political misinformation on a global scale. During the 2016 UK Brexit referendum it was found that political bots played a small but strategic role in shaping Twitter conversations. The family of hashtags associated with the argument for leaving the EU dominated, while less than one percent of sampled accounts generated almost a third of all the messages (Howard & Kollanyi, 2016).

False news reports, widely distributed over social media platforms, can in many cases be considered to be a form of computational propaganda. Bots are often key tools in propelling this disinformation across sites like Twitter, Facebook, Reddit, and beyond. These social media platforms have served significant volumes of fake, sensational, and other forms of junk news at sensitive political moments over the last several years. However, most platforms reveal little about how much of this content there is or what its impact on users may be. The World Economic Forum recently identified the rapid spread of misinformation online as among the top 10 perils to society (World Economic Forum, 2014). Prior research has found that social media favors sensationalist content, regardless of whether the content has been fact checked or is from a reliable source (Vicario et al., 2016). When junk news is backed by automation, either through dissemination algorithms that the platform operators cannot fully explain or through political bots that promote content in a preprogrammed way, political actors have a powerful set of tools for computational propaganda. Both state and non-state political actors can deliberately manipulate and amplify non-factual information online.

Cases and Methods

These case studies all begin with a basic set of research questions crafted for comparability. Is computational propaganda present in a country? What are its forms, types or styles? What is its impact on public life? Each case study also ends with some speculations. How might political bot activity run afoul of elections law in the country? Which computational propaganda campaigns had a significant impact, and how might they be prevented in the future?

These research findings are the result of knowledge generated through multiple social and data science methods. We have conducted qualitative and quantitative content analysis of news coverage about computational algorithm. We have done big data analysis of large networks of users on Facebook, Twitter and Weibo. Researchers used multiple methods in cataloguing their country-specific case studies including, but not limited to: interviews with users who have experienced attacks, interviews with those who have worked to produce political bots and social media-based propaganda and harassment, process tracing, participant observation, social network analysis, and content analysis of media articles. Each case required different approaches and tools. Researchers made use of both qualitative and

quantitative methods of analysis. This mixed-method approach enables the case studies to speak to concerns at the intersection of several disciplines: especially those focused on social science, law, and computer science.

The authors of these case studies were chosen for their knowledge of relevant languages and political cultures, their ability to conduct in-country interviews as needed, and their skills at analyzing large datasets of social media content where relevant. The research team involved 12 researchers across nine countries who, altogether, interviewed 65 experts, analyzed tens of millions posts on seven different social media platforms during scores of elections, political crises, and national security incidents.

Table 1: Evidence Used Across Country Case Studies

Country	Data Analysis	Interview Subjects	Platforms	Social Media and Politics
Brazil	281,441 tweets were collected from 82,575 unique accounts, collected February–March 2017, and 80,691 tweets from 33,406 unique users from in May 2017.	10	Facebook Twitter WhatsApp	Bot networks and other forms of computational propaganda were active in the 2014 presidential election, the constitutional crisis, and the impeachment process. Highly automated account support and attack political figures, debate issues such as corruption, and encourage protest movements.
Canada	3,001,493 tweets collected from September–October 2015.	10	Twitter	Political parties use bots, but there are also positive ways to use algorithms and automation to improve journalism and public knowledge.
China	1,177,758 tweets from 254,132 unique accounts, collected February–April on Twitter; 1,543,165 comments from 815,776 unique users on Weibo collected January–February 2017.	2	Facebook Twitter Weibo	On Twitter, several large bot networks published anti-government messages in simplified Chinese. Opinion manipulation on Weibo occurs, but not through automation.
Germany	121,582 tweets from 36,541 users, collected over the period of three days in February 2017; and 154,793 tweets from 32,008 unique users collected over the course of seven days in March 2017.	13	Facebook Twitter	Social bots played a marginal role in German elections; whereas substantial misinformation has been circulated during pivotal moments of political life. Germany has emerged as a leader in countering computational propaganda, with a state-wide regulation to be implemented in the summer, and numerous civil society watchdog projects.
Poland	50,058 tweets from 10,050 unique accounts collected March–April 2017 on Twitter.	10	Facebook Twitter	There is a clear industry of producing and managing fake accounts and automation over multiple platforms. A tiny number of right wing accounts generate 20% of the political content over Twitter.
Russia	14 million tweets collected from February 2014 to December 2015 from more than 1.3 million users.	0	Twitter	Russian Twitter networks are almost completely bounded by highly automated accounts, with a high degree of overall automation.
Taiwan	49,541 comments and replies to a message from the Taiwanese President in January 2016. 1,396 tweets about the President from 596 unique users collected in April 2017.	10	Facebook Twitter LINE	Combined human and automated personal and political attacks on the Taiwanese President.
Ukraine	Representative sample of political perspectives on MH17 tragedy, beginning summer 2014.	0	Facebook Odnoklassniki Twitter VKontakte	Ukraine is the frontline of experimentation in computational propaganda, with active campaigns of engagement between Russian botnets, Ukraine nationalist botnets, and botnets from civil society groups.
USA	17 million tweets from 1,798, 127 unique users, collected November 2016.	15	Facebook Twitter	Twitter bots constituted over 10% of the sample, and they reached highly influential network positions within retweet networks during the 2016 US election. The botnet associated with Trump-related hashtags was 3 times larger than the botnet associated with Clinton-related hashtags.

In Sum

We face new challenges in the investigation of automation and fake accounts on social media. First, we have found that political actors are adapting their automation in response to our research. This suggests that the campaigners behind fake accounts and the people doing their "patriotic programming" are aware of the negative coverage that this gets in the news media.

Second, we have found several kinds of bot networks that are quite active but that fall below our formal threshold of what counts as a bot. For countries where Twitter is not a particularly important social media platform, it seems that bots are prevalent but not performing as efficiently as bot networks in countries with lots of Twitter users. In many countries there are large numbers of "sleeper bots." These are accounts that have only tweeted a few times, usually in scattered ways, and have other account features that suggest automation.

Third, it is difficult to put research findings into service for public policy recommendations in consistent ways across countries because the legal questions about computational propaganda vary greatly from country to country. During the 2015 election in Canada, comedienne Sarah Silverman encouraged Canadians to vote for the National Democratic Party over Twitter. Is she a foreigner influencing voters in contravention of the Canada Elections Act? If bots propagate her message after campaigning is supposed to stop, are platforms or bot writers interfering with the election?

The advantage of cross-national comparisons is in yielding evidence about which policy responses can work well. In Taiwan, the government has responded with an aggressive media literacy campaign, and bots that will check facts for the public. In Ukraine, the government response has been minimal, but there are a growing number of private firms trying to make a business of fact checking and protecting social media users.

Automated political communication involves the creation, transmission, and controlled mutation of significant political symbols over expansive social networks. Indeed, the impact of digital information infrastructure on how political culture is produced is at least as interesting, though under-studied, as the impact of infrastructure on how political culture is consumed. While we can theorize about the

ways in which computational propaganda may violate political values or the social contract writ large. But the case studies in this collection of working papers demonstrate the origins and very concrete consequences of computational propaganda.

It is time for social media firms to design for democracy. For democracies, there are big elections ahead. Germany votes in late 2017. Egypt, Brazil, and Mexico all have general elections in 2018. In the US, strategists are already planning for the 2018 mid-term elections. Let's assume that authoritarian governments will continue to use social media as a tool for political control. But for democracies, we should assume that encouraging people to vote is a good thing. Promoting political news and information from reputable outlets is crucial. Ultimately, designing for democracy, in systematic ways, will help restore trust in social media systems.

Computational propaganda is now one of the most powerful tools against democracy. Social media firms may not be creating this nasty content, but they are the platform for it. They need to significantly redesign themselves if democracy is going to survive social media.

References

Dewey, C. (2016, November 17). Facebook Fake-News Writer: "I Think Donald Trump is in the White House Because of Me." *The Washington Post*. Retrieved from https://www.washingtonpost.com/news/the-intersect/wp/2016/11/17/facebook-fake-news-writer-i-think-donald-trump-is-in-the-white-house-because-of-me/?utm_term=.30dba5468d15

Forelle, M., Howard, P., Monroy-Hernández, A., & Savage, S. (2015). Political Bots and the Manipulation of Public Opinion in Venezuela. *arXiv:1507.07109 [Physics]*. Retrieved from http://arxiv.org/abs/1507.07109

Gallacher, J., Kaminska, M., Kollanyi, B., Yasseri, T., & Howard, P. N. (2017). Social Media and News Sources during the 2017 UK General Election. Retrieved from comprop.oii.ox.ac.uk

Howard, P. N. (2011). *Castells and the media*. New York, NY: Polity Press.

Howard, P. N., Bolsover, G., Kollanyi, B., Bradshaw, S., & Neudert, L.-M. (2017). Junk News and Bots during the U.S. Election: What Were Michigan Voters Sharing Over Twitter? Data Memo 2017.1. Oxford, UK: Project on Computational Propaganda. Retrieved from http://comprop.oii.ox.ac.uk/2017/03/26/junk-news-and-bots-during-the-u-s-election-what-were-michigan-voters-sharing-over-twitter/

Howard, P. N., & Kollanyi, B. (2016). Bots, #StrongerIn, and #Brexit: Computational Propaganda during the UK-EU Referendum. *arXiv:1606.06356 [Physics]*. Retrieved from http://arxiv.org/abs/1606.06356

Howard, P. N., Kollanyi, B., & Woolley, S. C. (2016). Bots and Automation Over Twitter during the US Election. *Computational Propaganda Project: Working Paper Series*.

Howard, P. N., & Parks, M. R. (2012). Social Media and Political Change: Capacity, Constraint, and Consequence. *Journal of Communication*, *62*(2), 359–362. https://doi.org/10.1111/j.1460-2466.2012.01626.x

Parkinson, H. J. (2016, November 14). Click and elect: how fake news helped Donald Trump win a real election. *The Guardian*. Retrieved from https://www.theguardian.com/commentisfree/2016/nov/14/fake-news-donald-trump-election-alt-right-social-media-tech-companies

Read, M. (2016, Npvembe). Donald Trump Won Because of Facebook. *New York Magazine*. Retrieved from http://nymag.com/selectall/2016/11/donald-trump-won-because-of-facebook.html

Vicario, M. D., Bessi, A., Zollo, F., Petroni, F., Scala, A., Caldarelli, G., … uattrociocchi, W. (2016). The Spreading of Misinformation Online. *Proceedings of the National Academy of Sciences*, *113*(3), 554–559. https://doi.org/10.1073/pnas.1517441113

Woolley, S. C., & Howard, P. N. (2016). Automation, Algorithms, and Politicsl Political Communication, Computational Propaganda, and Autonomous Agents — Introduction. *International Journal of Communication*, *10*(0), 9.

World Economic Forum. (2014). 10. The Rapid Spread of Misinformation Online. Retrieved March 8, 2017, from http://wef.ch/GJAfq6

About the Authors

Samuel C. Woolley is the Director of Research of the Computational Propaganda Project at the Oxford Internet Institute, University of Oxford, and a doctoral candidate (ABD) at the University of Washington. He is a research fellow at Google Jigsaw, the Institute for the Future, the TechPolicy Lab at the University of Washington and a former fellow at the Center for Media, Data and Society at Central European University. He researches automation and propaganda.

Philip N. Howard is a statutory Professor of Internet Studies at the Oxford Internet Institute and a Senior Fellow at Balliol College at the University of Oxford. He has published eight books and over 120 academic articles and public essays on information technology, international affairs and public life. Howard's books include The Managed Citizen (Cambridge, 2006), the Digital Origins of Dictatorship and Democracy (Oxford, 2010), and most recently, Pax Technica: How the Internet of Things May Set Us Free or Lock Us Up (Yale, 2015). He blogs at www.philhoward.org and tweets from @pnhoward.

Citation

Series Acknowledgements