National Endowment for Democracy
Supporting freedom around the world

forum
International Forum for Democratic Studies

FORUM Q&A: PHILIP HOWARD ON COMPUTATIONAL PROPAGANDA'S CHALLENGE TO DEMOCRACY

Philip N. Howard (@pnhoward) is the principal investigator of the Computational Propaganda Project at the Oxford Internet Institute at the University of Oxford. A professor and writer, Howard has authored numerous academic articles, essays and books on information technology, international affairs, and public life, as well as on the use of digital media for both civic engagement and social control in countries around the world. His projects on digital activism, information access, and modern governance in both democracies and authoritarian regimes have been supported by the National Science Foundation, U.S. Institute of Peace, and Intel's People and Practices Group. His most recent book is Pax Technica: How the Internet of Things May Set Us Free or Lock Us Up.

Computational propaganda, or the use of algorithms and automated social media accounts to influence politics and the flow of information, is an emerging challenge to democracy in the digital age. Using automated social media accounts called bots (or, when networked, botnets), a wide array of actors including authoritarian governments and terrorist organizations are able to manipulate public opinion by amplifying or repressing different forms of political content, disinformation, and hate speech.

Dean Jackson of the International Forum for Democratic Studies spoke with Phil Howard to discuss political bots, computational propaganda, and the challenges they pose to democracy. (This interview has been condensed and edited for clarity. The views and opinions expressed here do not necessarily reflect those of the National Endowment for Democracy.)

Dean Jackson: You call your project the "Computational Propaganda Project." Propaganda is usually seen as an attempt to persuade. What are some of the ways in which computational propaganda has been used to persuade human consumers and influence political discussions?

Phil Howard: There have been some real, concrete examples of how computational propaganda has been used to persuade. There were multiple algorithms operating over Twitter and multiple 'fake news' stories over Facebook over Hillary Clinton's alleged corruption and involvement with a pedophilia ring and the suicide of an FBI agent, multiple news stories that we know now are fake news. But because of the way that social media transmits information, significant numbers of American voters still think that she was involved in a pedophilia ring in a pizza parlor. So you can tell that some of these outcomes in public opinion are largely the result of coordinated misinformation campaigns.

It's often remarked that propaganda and censorship are two sides of the same coin. It seems like this is true of computational propaganda, as well. How can bots and botnets be used to suppress information as well as propagate it?

That's a very good question. Some of the earliest uses of automated accounts, often called 'bots,' were to suppress public conversations over what was going on in Syria. There was a moment when the Syrian civil war broke out in 2010 when very few news outlets had journalists stationed in Damascus. We were all relying on Twitter for photos and stories and vignettes about what was going on in the civil war. It is believed that the Al Assad government hired a communications firm based in Bahrain to flood the Syria hashtag with soap opera stories, scores from soccer games, and pictures of the beach, effectively choking off that hashtag as a useful way of learning about the crisis. So, for the most part, that kind of automation can help drive people off a hashtag that's important in a crisis.

One of the reasons this question is important is that a number of governments are working on ways to manage political conversation, ways that lead to overregulation and that may also generate their own forms of censorship. So what we need is some sort of careful policy path between firms that serve junk to users, on the one hand, and governments that overregulate.

Only a few years ago, most observers would have guessed that bots were mostly ignored by human users. People seemed to believe that we could easily identify automated activity, and so we ignored it or found it merely obnoxious or annoying. But your research suggests this isn't the case today, if it ever was. Is the sense that bots have grown more sophisticated over time correct, and if so, in what ways have they done so?

They've grown more sophisticated in several ways. Even back then, way back in 2010 or 2011, bots may have been used for generating junk, which we can identify, but they were also used to make politicians seem more popular than they actually were. So even the bots that weren't generating content helped boost your follower numbers and made it look like you were more popular than you actually were, and like there was a community of support behind you.

These days, bots are more sophisticated in two ways: first, they're often loaded with many different complex messages that sometimes lead humans into semi-serious, engaged conversations. Some of our research in our latest report on political activity in the United States actually shows that we can measure how bots move from being peripheral—many bots are peripheral to social networks—to being central with what you call a K-Core measure. So you can actually identify highly automated accounts that become, over time, very popular among humans.

The other thing is that many of these accounts are also tied to fairly complete profiles that include Facebook profiles, SIM cards, Gmail addresses—a whole range of different technologies—that make these fake accounts seem like real users, real voters.

Some observers, including Farhad Manjoo at *The New York Times*, have commented on bots' ability to influence users who set the news agenda, such as journalists. Are there other ways in which bots have influence? Can bots instead capitalize on existing social debate and divide to influence more typical and average users?

Usually, that kind of influence comes through negative campaigning, like messaging around Hillary Clinton and corruption, or saying no to a new policy idea. It's very rare that you see an effective automated communications campaign about an exciting, new, positive policy idea. You just don't see those very much. It's the negative ones, the ones that are angry, the ones that say "no," or "we've had enough," the ones that express outrage, the ones that pass moral judgement; it's usually the ones that come with some picture of a politician at an odd angle that makes them look like they're dark and scary… that kind of messaging is what tends to travel the farthest amongst networks of average users.

I want to ask about networks of grassroots users, especially groups of human users who coordinate trolling and disinformation campaigns on their own. This has been a particular issue around national elections, including in France and the United States, and now also in Germany's upcoming election. How much interaction do these grassroots campaigns have with automated campaigns? Do they link up at all?

We don't know for sure whether they "link up" organizationally. We can identify different networks that have generated content, but we don't know that the people behind those networks are communicating. What we know is that there's a network of Russia-based accounts that follow a number of U.S. politicians, share content, share news from Russia Today and Sputnik. We know that [U.S. President Donald] Trump has many bot followers. There is a Trump fan base that generates a lot of content, and then there's sort of an alt-right network of content that's produced in the United States by bot-writers who live in [places like] Seattle, San Francisco, Brooklyn, and Montana, who aren't exceptional except in their ideology.

We don't know that these groups communicate with each other organizationally, but what we know is that they generate remarkably similar content, and sometimes they pick up each other's content. But, these days, it's getting harder and harder to separate networks.

It seems like most of this research focuses on Twitter, because that platform is easier to both manipulate and research. Are bots also influential on Facebook or on direct-messaging apps that are more commonly used as news sources in the developing world?

Unfortunately, we're having trouble studying other platforms. Twitter's Application Programming Interface (API) is the most accessible, and that's why we spend the most time on Twitter. Facebook's API is less accessible, they don't share data, they don't work well with academics, they don't really respond to queries, and they never provide replication data. We know that they hire many data scientists who conduct experiments internally, but we assume those are used to improve the delivery of advertising or the user experience and are not designed to improve public life.

In some of these newer direct-messaging platforms, they're too specialized… some of our teams have been able to study Line, which is used in Taiwan a lot. Some of these platforms are specific to Southeast Asia, and of course China has its own suite of social media platforms that are different again.

I'm curious, just as a follow-up: when you are able to study those direct-messaging apps, do you find a prevalence of bot activity, or is it still too hard to tell?

I'd say it's too hard to tell. In the next year, we're going to look at the use of direct-messaging apps in Iran.

That sounds like a really interesting area of future research, and a good segue: in your opinion, what are the most promising paths forward for your research? What issues most urgently demand the attention of researchers focusing on political bots and computational propaganda?

The most urgent issue is more of a policy one than anything: I think it's about getting Facebook to share. Most of what we know about social networks and political opinion formation is through Twitter data, and all of us who are researching this just kind of hope or assume that we can generalize from Twitter to Facebook networks. There are some good reasons to do that, but we know that there are entire countries where public life is on Facebook, not on Twitter. I think the most urgent thing is for academics to engage with Facebook, or for policymakers to intervene or work out ways that Facebook can help us tackle some of the big misinformation campaigns around the world.

You've indicated that the problem is about misinformation spread by automation. To what degree is the problem really about cognition, human behavior, the way we interact with information we encounter on the internet, and the way bots are able to game that? Do we have a good understanding of how this problem plays out inside of our heads, rather than inside of our devices?

It's a great point; I believe that there are several kinds of cognitive explanations for what goes on that are totally plausible. There are several selective exposure arguments: one is that we don't like to be contradicted; one is that once we've made a decision, we come up with shortcuts for what we feel we've already learned; and one is that we just choose to hear good things from the people we always spend time with. Those are three variations on the "selective exposure argument."

I think there are several levels of answers to the problem. The big picture answer is civic education. Teaching everybody Aristotle's top argumentative fallacies in Latin and how to spot them, so that everyone would leave high school knowing what an *ad hominem* attack is, knowing what an *argumentum ad populum* is—you probably could learn the Latin, but—just being able to identify ten argumentative fallacies and to understand what they are would be a great achievement, and it would help a lot. But that is super macro, to get every American to know what logical fallacies are.

The most proximate cause of the problem is that misinformation is presented on social media more frequently in the days before people vote. We've actually been able to demonstrate that the proportion of professionally produced news to junk news is at an all-time low the night before an election. We found this in Michigan in our Michigan sub-sample; we found it again in the UK during the recent UK election.

My last question is about the effects of near-future advances in technology, and how they might make the challenges of computational propaganda much more acute: for example, forms of video-editing that can make it appear like someone said something they did not on live television, or the democratization of other tools that allow for more sophisticated forms of misinformation. What is going to get worse most quickly if we don't address this challenge?

I would look more to the back-end than to video-editing, which could be a problem, but the two big back-end issues will be the application of real machine-learning and artificial-intelligence algorithms to political communication. Once the algorithms get genuinely interactive enough that you can pre-load content and engage people who really can't tell that they're engaging a bot, that will be a significant problem for political communication.

The second back-end issue involves the incorporation of new data from the internet of things. Much of the data that currently drives political messaging is credit card data, merged with health data, merged with voter registration files. Once there is additional behavioral data from lightbulbs, and your car, and the chip in your phone, and all of the data from the internet of things becomes the food for computational propaganda, that will make a seriously deep set of knowledge for political engagement with particular citizens. Those are the two back-end things that I think are the most concerning.