

Political Communication, Computational Propaganda, and Autonomous Agents

Introduction

SAMUEL C. WOOLLEY
University of Washington, USA

PHILIP N. HOWARD
University of Oxford, UK

The Internet certainly disrupted our understanding of what communication can be, who does it, how, and to what effect. What constitutes the Internet has always been an evolving suite of technologies and a dynamic set of social norms, rules, and patterns of use. But the shape and character of digital communications are shifting again—the browser is no longer the primary means by which most people encounter information infrastructure. The bulk of digital communications are no longer between people but between devices, about people, over the Internet of things. Political actors make use of technological proxies in the form of proprietary algorithms and semiautomated social actors—political bots—in subtle attempts to manipulate public opinion. These tools are scaffolding for human control, but the way they work to afford such control over interaction and organization can be unpredictable, even to those who build them. So to understand contemporary political communication—and modern communication broadly—we must now investigate the politics of algorithms and automation.

Keywords: political communication, science and technology studies, research trends, bots, algorithms, automation, social media, Internet of things

The vast majority of political speech acts now occur over digital platforms governed by terms-of-service agreements. In volumes of data or proportions of bandwidth, most communication is between and about devices and about people. It used to be fairly straightforward to trace agency—or to place blame—for miscommunication or communication that promulgated social inequality. Now significant amounts of communication involve autonomous agents that have been purposefully designed but that only produce content in interaction with people in the context of a platform.

Samuel C. Woolley: samwooll@uw.edu

Philip N. Howard: philip.howard@oii.ox.ac.uk

Date submitted: 2016–09–03

Copyright © 2016 (Samuel C. Woolley & Philip N. Howard). Licensed under the Creative Commons Attribution Non-commercial No Derivatives (by-nc-nd). Available at <http://ijoc.org>.

The code behind social media services, from Facebook to WhatsApp, often obfuscates the intentions of those who write it. Indeed, iterative construction of computer scripts—and the networked nature of digital authorship—mean that the outcomes of an algorithm can be unpredictable to both designers and publics. Politically motivated software agents are also deployed on the front end of platforms in ways that extend the same complications to the procedures that govern these important communication channels. Autonomous programs are used as proxies for political actors hoping to sway public opinion through the spread of propaganda and misinformation. Researching how algorithms and automation structure our lives is the next great challenge for the social sciences because such research involves a bewildering array of actors, artifacts, and code in complex webs of causality.

The software engineers, computer scientists, hackers, and sophisticated technology users who automate communication on social or device networks often call their creations bots. The word *bot* is an abbreviation of *robot*, itself a 20th century Czech term meaning “forced labor” or “slave.” This is a fitting etymology for bots in that they exist, in a manner of speaking, as digital versions of their embodied cousins. The majority of bots function as computationally enhanced conduits for human coders: They scrape the Web for information, monitor chat sites for misuse, and crawl financial databases for market trends. Most of the bots discussed in this Special Section are unique, however, in that they are not only interactive but also politically oriented.

The articles included in this Special Section were developed for, and through, a workshop held before the 2016 International Communication Association meetings in Fukuoka, Japan. The call for papers was issued in December 2015. We received 30 submissions; five were rejected on the basis of fit, and 25 were reviewed by the editors. In January 2016, 15 of these were selected for the preconference, and the editors provided feedback to authors. The papers were revised and shared with other participants in May 2015, and the workshop was held on June 8. Each paper had a dedicated respondent, and during each presentation a rapporteur recorded group questions and feedback for each author. In addition, the editors collected other manuscripts presented during the main conference and evaluated them for topical fit and analytical rigor. After the conference, authors of the most advanced manuscripts were invited to make revisions and be considered for inclusion in this collection of papers. The Special Section editors—and the editing team at *The International Journal of Communication*—provided an additional round of written review on the final collection of accepted manuscripts.

The workshop was supported by the European Research Council–funded Computational Propaganda (COMPROP) project of the Oxford Internet Institute (politicalbots.org). Any opinions, findings, conclusions, or recommendations expressed in this material are those of the manuscript authors and do not necessarily reflect the views of the European Research Council.

The Changing Demography of Users and Devices

Several big technology diffusion trends explain why algorithms, automation, social media, and the Internet of things are emerging as they are (Howard, 2015). First, device networks are spreading quickly. Our planet now holds slightly more than 7 billion people, it is home to more than 10 billion mobile

digital devices (Cisco, 2014). The number of stationary—yet connected—devices is estimated at 40 billion. This is both an immense network of opportunity and a significant infrastructure for abuse. It is also a growing network. The world will have more wired, intelligent devices in the coming years. The network is already made up of more devices than people, and the devices continue to share information even when we are not personally using them.

Second, more people are actively or passively generating data that is captured and transported over the Internet. By most estimates, in 2020 everyone will effectively be online. Most of the world will have direct Internet access through mobile phones and the Internet, but everyone will be immersed in a world of devices that are constantly connected to the Internet. People who lack direct access are tracked and monitored through government and corporate databases, and their economic, political, and cultural opportunities are still shaped by digital media. This means that for the first time in history, virtually everybody can connect to virtually everyone else. Most countries have upwards of 80% Internet penetration. And this is only the Internet of mobile phones and computers. The Internet of things will keep everyone networked constantly.

Third, most Internet users—and eventually most people—will soon be digital natives. Digital natives are the people born since the turn of the century in countries where device networks are a ubiquitous part of social life. By 2010, the majority of all Internet users were digital natives. Until recently, those using the Web were largely a kind of digital immigrant population that had started going online when they arrived at their universities or had become connected through their work. This is changing. Now more than 50% of the Internet-using population consists of people who were born into a world of pervasive Internet, mobile phones, and digital media. Alas, this population is unevenly distributed, and most of these technology users reside in wealthy countries. These days, however, the bulk of new technology users are young and living in Africa or Asia. Most are coming from failed, fragile states with weak economies.

The user norms that shaped the Internet of the late 1990s and early 2000s will be superseded by the practices and values of the next 2 billion users. What will the Internet look like after young people in Tehran, Nairobi, and Guangzhou reshape its content? How will political communication scholars prepare to study the mediation that occurs over such complex networks, with such a variety of human, quasihuman, parahuman, and pseudohuman agents, actants, subjects, and audiences?

The Changing Agency Behind Political Communication

Social media has had a profound impact on the way people discuss current affairs and engage with politics. Presence on these platforms helps young people independently cultivate a political identity and provides users of all ages with personalized civic knowledge in both authoritarian and democratic regimes. Large-scale grassroots movements have been born, organized, and disseminated on platforms such as Facebook, YouTube, WhatsApp, FireChat, and Twitter. The positive, and wide-reaching, democratic potential of social media is much discussed, but another, more propagandistic, side of this new technology exists (Woolley & Howard, 2016a).

Bots—algorithmically driven computer programs designed to do specific tasks online—have invaded political conversations worldwide. Bots make up nearly 50% of all online traffic and account for a significant portion of active users on the most popular social media platforms (Zeifman, 2015). On Twitter, for instance, approximately 30 million active accounts are bot driven—they mimic human users and produce copious information (Motti, 2014).

The pervasive use of such human-software hybrids, and the obscure and often discriminatory nature of the algorithms behind them, threaten to undermine the political potential—organizational, communicative, and otherwise—of social media systems (boyd, Levy, & Marwick, 2014; Woolley, 2016; Woolley & Howard, 2016b;). Many types of actor groups build, use, and deploy political bots: corporate lobbyists, content management firms, civic activists, defense contractors, and political campaigns. This Special Section explores the rise of political bots through platform- and country-specific case studies, automation-oriented policy analyses, and ethical interrogations of artificially intelligent systems.

Four Definitions

Every new line of inquiry comes with particular keywords and contested definitions. In our own work on political algorithms, and looking across the articles in this Special Section, we find four particular terms worth explaining. We have borrowed from traditional definitions and reworked them in light of new evidence. Our goal here is to explain how those of us working on algorithms, automation, and politics are both playing with conventional keywords and creating new ones: social media, political bots, the Internet of things, and computational propaganda.

The contributors to this Special Section have an expansive understanding of what social media is, and they use the term in at least one of the ways defined by Howard and Parks (2012). *Social media* consists of (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 become cultural products; and (c) the people, organizations, and industries that produce and consume both the tools and the content (Howard, 2011; Howard & Parks, 2012). Consequently, *political bots* are the algorithms that operate over social media, written to learn from and mimic real people so as to manipulate public opinion across a diverse range of social media and device networks. Such bots are a variety of automated computer scripts that interact with other users on social media platforms such as Twitter and community-maintained sites such as Wikipedia. Political bots are deployed, for example, to boost follower numbers and to retweet the content of political candidates on Twitter, to attack political opponents on Facebook, or to drown out activists' conversations on Reddit.

Not all of the political algorithms worth investigating operate over social media, but that is where the most overt and aggressive of them work. Some operate over the Internet of things, which is rapidly turning everyday objects into media and necessitating that communication scholars expand their purview. The Internet of things is defined as the networks of manufactured goods with embedded power supplies, small sensors, and addresses on the Internet. Most of these networked devices are everyday items that are sending and receiving data about their conditions and our behavior (Howard, 2015).

Altogether, social media, political bots, and the Internet of things enable computational propaganda. We define *computational propaganda* as the assemblage of social media platforms, autonomous agents, and big data tasked with the manipulation of public opinion. Autonomous agents, equipped with big data about our behavior collected from the Internet of things, work over social media to engage with us on political issues and advance ideological projects. Computational propaganda involves software programs that are interactive and ideologically imbued. They are interactive within the context of a platform. They are ideological, first, in that they are programmed to promote a particular perspective in politically charged conversations and, second, in that they are artifactual evidence of the idea that technology can be used to influence politics. These are almost pure examples of politics in code. Computational propaganda is among the latest, and most ubiquitous, technical strategies to be deployed by those who wish to use information technology for social control.

Contemporary Research on Political Bots, Algorithms, and Automation

Computer scientists and policy makers often treat bot-generated traffic as a nuisance to be managed, and system administrators work to simply shut down accounts that run automatically generated scripts. We disagree with this easy dismissal and argue that computational propaganda is becoming one of the most socially significant impacts of innovations in computer science and engineering. Developing good ways of understanding the innovation networks of bot makers and building new knowledge about politicized algorithmic systems is crucial for sustaining the public value of social media.

The articles in this collection contribute to the foundation for this new domain of research. Through tough ethical questions and empirical analysis, our authors assemble a series of suggestions for how to both theorize about political automation and respond to bot-driven propaganda. Our authors address the rise of the political bot from several substantive angles using varied research methods.

Within this Special Section, Guilbeault (2016); Maréchal (2016); Mittelstadt (2016); and Sandvig, Hamilton, Karahalios, and Langbort (2016) interrogate the ethics and regulation of algorithms and bots. Whereas Maréchal argues for normalized, state-based oversight of bias at the algorithmic level, Mittelstadt suggests that the burden of auditing these systems for political bias lies on the shoulders of the platforms themselves. Guilbeault develops an ecologically contextualized account of agency over digital networks, arguing that political bots provide a window for both theoretical and policy-oriented discussions about aspects of both innovation and intervention in the digital sphere. Sandvig and his collaborators demonstrate that the algorithm itself can be audited for ethical implications, and they convincingly argue that this will be an increasingly important research activity for the social sciences.

Neff and Nagy (2016) conduct a case study of Tay, Microsoft's now infamous Twitter bot. They ask, What is agency in a world occupied by both humans and bots? By reviewing the history of how chat bots have been built, they reveal limitations to our current understanding of human and algorithmic agency. They demonstrate that we need a perspective of "symbiotic agency" (Nagy & Neff, 2015). Ford, Puschmann, and Dubois (2016) explore the properties and impact of the "transparency bot" through an

examination of bots that track government employees' edits to Wikipedia. Although popular opinion suggests that these "WikiEdits bots" provide a neutral view of governments' Wikipedia activity, Ford et al. find that they negatively skew government contributions to the free online encyclopedia.

Through another exciting research model, Murthy et al. (2016) find that for political bot campaigns to be successful in manipulating opinion, they must have a significant amount of social, temporal, and monetary capital behind them. These findings come from an experiment in which the authors launched and tracked the progress of several accounts, including some supported by political bots, on Twitter. Kollanyi's (2016) article is one of the first studies of GitHub code production and the bot-production community there. It seeks to build systematic knowledge around bot production by analyzing available bot code on the open-source site.

Finally, Shorey and Howard (2016) provide a bookend to this Special Section with an analysis of the existing domains of inquiry into algorithms, automation, and politics. Critical studies like these are particularly important at the early stages of inquiry, when it is important to craft questions and understand the implications of new phenomena. Based on both a literature review and extensive consultations with the researchers working on these themes, this article annotates current inquiry, identifies the gaps in current inquiry, and suggests directions for future research.

Conclusion

In this Special Section—a collection of articles resulting from a preconference at the 2016 International Communication Association meetings—we find that algorithms govern the burgeoning communications between us. Such algorithms mediate almost all interaction and content that we do not experience directly, face-to-face and in person. We find people communicating, sometimes unawares, with automated scripts. Algorithms and automation invariably involve political choices in design, the data collected about our interactions invariably has political applications, and a growing number of people are targeted by computational propaganda.

Communication now occurs over a plethora of devices and platforms that we have built but do not fully control. We can identify the particular firms and engineers behind specific platforms and applications, but even those software engineers have trouble predicting what their algorithms will do when released to interact with users on a platform.

There are two additional methodological attributes to the current tone of this new domain of inquiry. First, the contributors adopt and adapt some of the analytical frames used in science and technology studies for the study of political communication. Analytically, this means being open to locating agency, if needed, with material objects and code, not just people. Instead of positivist insistence on the hard determinism test that technology provide a clear cause of some social effect, we prefer the soft determinism that allows the capacities and constraints of technology to provide boundaries for social action. Political speech acts that occur over social media, and large amounts of data collected over the Internet of things, shape what autonomous agents do. Instead of forcefully identifying causes with no a

priori context, most of the authors in this collection treat technology design and use as the inciting incident in their analytical narratives.

Second, the researchers behind these studies employ various forms of network ethnography. This method involves observing designers in their natural settings, treating software and hardware as artifacts, and tracing the impact of the designers and their artifacts through to a population of users (Howard, 2002). This sort of mixed-method approach—pairing participant observation with large-scale data analysis—can productively improve our understanding of how designers, artifacts, and users all fit together. Such methodological pairings are likely to become the standard in scholarly research on political communication in the years to come. Traditional training in correlational statistics was most useful for political communication scholars who thought that using public opinion data was the best route to generalization. With the growing complexity in information infrastructure and the burgeoning volumes of behavioral data, regression models will need to take a backseat to the more powerful combination of ethnographic and computational social science.

Whether or not traditional scholars of political communication jump into the analysis of autonomous agents and computational propaganda, it is very likely that the next generation of scholars will see this as an exciting domain of inquiry with direct implications for the quality of public life in the years to come. But whether we can actually *do* this research is an open question.

In the United States, the Computer Fraud and Abuse Act has associated many hacking activities with criminal endeavors, and the communication scholars who currently audit computer algorithms and run experiments to see how autonomous agents shape public life have recently sued the government. Dr. Christian Sandvig, a professor of communication at the University of Michigan and a contributor to this Special Section, asserts that his experiments run by opening fake profiles to pose as job and housing seekers constitutes speech and expressive activity that is protected under the First Amendment. He and his colleagues argue that future academic inquiry can be seriously hampered by the businesses that operate platforms over which so much communication occurs if those businesses dislike the research questions being asked. Such social media sites and proprietary device networks can change their terms of service at any time without informing visitors, turning any speech or activity on the site into a criminal act.

For Sandvig and others, this is a violation of the Fifth Amendment right to due process, which requires proper notice to the public of what constitutes criminal behavior (Farivar, 2016; Zetter, 2016). Increasingly, social media and digital devices generate the evidence that political communication scholars need to work with (Sandvig, Hamilton, Karahalios, & Langbort, 2014). As researchers, we need to attend to the impact of algorithms and automation on public life. We also need to assert our right of access to evidence about how public life is changing. This Special Section is a start to the scholarly conversation about algorithms and political communication, and a call to other scholars to join our efforts.

References

- boyd, d., Levy, K., & Marwick, A. (2014). The networked nature of algorithmic discrimination. In S. P. Gangadharan, V. Eubanks, & S. Barocas (Eds.), *Data and discrimination: Collected essays* (n.p.). Washington, DC: Open Technology Institute and the New America Foundation.
- Cisco. (2014). Cisco visual networking index: Global mobile data traffic corecast update, 2013–2018. *Cisco.com*. San Jose, CA: Cisco. Retrieved from http://cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/white_paper_c11-520862.html
- Farivar, C. (2016, June 29). To study possibly racist algorithms, professors have to sue the US. *Arstechnica.com*. Retrieved from <http://arstechnica.com/tech-policy/2016/06/do-housing-jobs-sites-have-racist-algorithms-academics-sue-to-find-out/>
- Ford, H., Puschmann, C., & Dubois, D. (2016). Keeping Ottawa honest, one tweet at a time? Politicians, journalists, Wikipedians and their Twitter bots. *International Journal of Communication, 10* (Special Section).
- Guilbeault, D. (2016). Political bots as ecological agents: The ethical implications of digital space. *International Journal of Communication, 10* (Special Section).
- Howard, P. N. (2002). Network ethnography and the hypermedia organization: New media, new organizations, new methods. *New Media & Society, 4*(4), 550–574. <http://doi.org/10.1177/146144402321466813>
- Howard, P. N. (2011). *Castells and the media*. New York, NY: Polity Press.
- Howard, P. N. (2015). *Pax technica: How the Internet of things may set us free or lock us up*. New Haven, CT: Yale University Press.
- Howard, P. N., & Parks, M. R. (2012). Social media and political change: Capacity, constraint, and consequence. *Journal of Communication, 62*(2), 359–362. <http://doi.org/10.1111/j.1460-2466.2012.01626.x>
- Kollanyi, B. (2016). Where do bots come from? An analysis of bot codes shared on GitHub. *International Journal of Communication, 10* (Special Section).
- Maréchal, N. (2016). When bots tweet: Toward a normative framework for bots on social networking sites. *International Journal of Communication, 10* (Special Section).
- Mittelstadt, B. (2016). Auditing for transparency in content personalization systems. *International Journal of Communication, 10* (Special Section).

- Motti, J. (2014, August 12). Twitter acknowledges 23 million active users are actually bots. Retrieved from <http://www.techtimes.com/articles/12840/20140812/twitter-acknowledges-14-percent-users-bots-5-percent-spam-bots.htm>
- Murthy, D., Powell, A., Tinati, R., Anstead, N., Carr, L., Halford, S., & Weal, M. (2016). Bots and political influence: A sociotechnical investigation of social network capital. *International Journal of Communication, 10* (Special Section).
- Nagy, P., & Neff, G. (2015). Imagined affordance: Reconstructing a keyword for communication theory. *Social Media + Society, 1*(2). <http://doi.org/10.1177/2056305115603385>
- Neff, G., & Nagy, P. (2016). Talking to bots: Symbiotic agency and the case of Tay. *International Journal of Communication, 10* (Special Section).
- Sandvig, C., Hamilton, K., Karahalios, K., & Langbort, C. (2014). Auditing algorithms: Research methods for detecting discrimination on Internet platforms. *Social.cs.uiuc.edu*. Retrieved from <http://social.cs.uiuc.edu/papers/pdfs/ICA2014-Sandvig.pdf>
- Sandvig, C., Hamilton, K., Karahalios, K., & Langbort, C. (2016). When the algorithm itself is a racist: Diagnosing ethical harm in the basic components of software. *International Journal of Communication, 10* (Special Section).
- Shorey, S., & Howard, P. N. (2016). Automation, big data, and politics: A research review. *International Journal of Communication, 10* (Special Section).
- Woolley, S. (2016). Automating power: Social bot interference in global politics. *First Monday, 21*(4). <http://doi.org/10.5210/fm.v21i4.6161>
- Woolley, S., & Howard, P. N. (2016a). Social media, revolution, and the rise of the political bot. In P. Robinson, P. Seib, & R. Frohlich (Eds.), *Routledge handbook of media, conflict, and security* (pp. 282–292). New York, NY: Routledge.
- Woolley, S., & Howard, P. N. (2016b, May 15). Bots unite to automate the presidential election. *Wired*. Retrieved from <http://www.wired.com/2016/05/twitterbots-2/>
- Zeifman, I. (2015, December 9). 2015 bot traffic report: Humans take back the Web, bad bots not giving any ground. *Incapsula.com*. Retrieved from <https://www.incapsula.com/blog/bot-traffic-report-2015.html>
- Zetter, K. (2016, June 29). Researchers sue the government over computer hacking law. *Wired*. Retrieved from <https://www.wired.com/2016/06/researchers-sue-government-computer-hacking-law/>