Manufacturing Consensus:
Computational Propaganda and the 2016 U.S. Presidential Election

Samuel Christopher Woolley

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington
2018

Reading Committee:
Philip N. Howard, Chair
Benjamin Mako Hill
Gina Neff

Program Authorized to Offer Degree:
Department of Communication

University of Washington

Abstract

Manufacturing Consensus:
Computational Propaganda and the 2016 U.S. Presidential Election

Samuel Christopher Woolley

Chair of Supervisory Committee
Dr. Philip N. Howard

This dissertation is an investigation of the ways three political actor groups used and interacted with bots and computational propaganda during the 2016 Presidential Election in the United States of America: political campaigns, journalists, and digital constituents. It is informed by data from over nine months of fieldwork, from February 2016 to after November 2016, in and around the U.S. campaign including attendance at pivotal party events, participant observation of campaigns, and consistent collection and parsing of online and offline information related to bots, computational propaganda, and the race. Over 40 interviews with a variety of experts from each of the three actor groups were done for this project. Three core working theoretical concepts emerged from this research: manufacturing consensus, the bot as an information radiator, and the bot as a proxy for the creator. The first describes the usage of bots and computational propaganda in attempts to amplify content online and give political ideas and actors the illusion of popularity in an effort to create bandwagon support. The second related to how journalists use bots as prostheses for reporting—to write simple stories, collect and parse data, and continuously communicate with the public about important information. The third explores how digital constituents, or citizens engaged in digital political communication and the usage of bots and computational propaganda online, and other bot builders can be theoretically conceived as related to, but separate from, the bots they create and deploy over the Internet. I argue, via each

of these concepts and through the other findings of this research, that bots are one of the most important new tools for political communication in the United States. These automated software actors are also useful in understanding novel relationships between technology and society.

Table of Contents

List of Figures

Glossary

API: Application Programming Interface

DDoS: Distributed Denial of Service

DNC: Democratic National Committee

IFTTT: If This Then That

GOP: Grand Old Party

GOTV: Get Out the Vote

MAGA: Make America Great Again

NDI: National Democratic Institute

NED: National Endowment for Democracy

NGO: Non-Governmental Organization

PAC: Political Action Committee

RNC: Republican National Committee

Acknowledgements

For his steady support, guidance and feedback, I would like to acknowledge my chair Dr. Philip Howard. I would also like to offer earnest thanks to Benjamin Mako Hill, Gina Neff, Ryan Calo, and the Computational Propaganda Research Project team at the University of Oxford for their generous help. I am grateful to the faculty at the Department of Communication at the University of Washington and the Oxford Internet Institute at the University of Oxford. Research collaborations and friendships with fellow scholars Douglas Guilbeault, Nick Monaco, Samantha Shorey, and Tim Hwang were integral to my development as a researcher of bots and computational propaganda. I am also thankful for fellowships at Google's think-tank Jigsaw, Central European University's Center for Media, Data and Society, the Tech Policy Lab at the University of Washington and the Institute for the Future, which stretched my thinking about the political use of automated technology. I developed ideas in portions of this dissertation through feedback from several workshops, talks and paper presentations at: the Stanford Hoover Institution & Freeman Spogli Institute, the Yale Information Society Project, the Data and Society Research Institute, the Annenberg-Oxford Media Policy Summer Institute, and the Princeton Center for Information Technology Policy. The collegiality and thoughtfulness of summer programs including the Annenberg-Oxford Media Policy Summer Institute and the CEU-Annenberg course of Advocacy, Activism, and the Internet contributed to my growth as a scholar of internet studies. Please direct any correspondence to Box 353740, Communication Department, University of Washington, Seattle, WA 98105. Parts of the research for this project were supported by the National

Science Foundation, the Ford Foundation and the European Research Council. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the University of Washington, the National Science Foundation, the Ford Foundation or the European Research Council.

Dedication

To my wife, Samantha.

Chapter 1

Introduction: The Political Bot

A. Introduction: The Digital Wild West

A week before the 2016 United States presidential election I received a call from a contact at the National Democratic Institute (NDI), a non-profit organization loosely affiliated with the U.S. Democratic Party. He informed me that representatives from the U.S. State Department, several prominent D.C. based NGOs, two or three major U.S. media outlets, and some of the major U.S. technology firms were hosting a small invite-only convening in Silicon Valley two days after the election. The plan was to do a post-hoc analysis of the role of social media, automation, and artificial intelligence during the 2016 contest. My contact at NDI told me that had been following my work on the use of social media bots, computational propaganda, and U.S. politics and wanted me to give a keynote talk on the subject. I accepted the invitation. At that stage, Clinton's lead was "holding steady" in most of the major polls and my contact seemed confident about the outcome of the hard-fought election (Lauter, 2016).

On November 9th, 2016, the political landscape of the U.S. had changed markedly. Along with this came a cascade of questions about the efficacy of traditional campaign political communication strategy during U.S. elections. Not only had Donald J. Trump won, but he'd done so much to the shock of most of the journalists and digital communication experts I had gotten to know over nine months of fieldwork preceding election day. The Republican candidate and his campaign used novel political communication tactics in

their path to victory. They prioritized social media advertising above any and all traditional media. In fact, "social", as the consultants called it, was the only media marketing area the Trump campaign outspent Clinton and her campaign (Wesleyan Media Project, 2017). Throughout 2016 reporters and academics had observed that online communication during pivotal electoral events, especially information on major platforms such as Facebook and Twitter, had been consistently manipulated by armies of automated bots (Andrews, 2016; Howard et al., 2016; Bessi & Ferrara, 2016a). Concern about the role of digital disinformation—or fake news—was burgeoning amongst pundits.

By November 10th, 2016, the day I stood in front of the group of confused D.C. and Silicon Valley insiders, I had spent hundreds of hours talking to people who built, used, tracked, and interacted with political bots during the contest. Political bots are best defined as software programs used to automate social media profiles in attempts to manipulate public opinion (Woolley & Howard, 2016a). I had talked to digital consultants who had worked for both major U.S. political parties, retired computer programmers who had found new ways to game online trends, and amateur journalists who used bots to drive traffic to their stories about the election. I had also spoken to reporters at flagship U.S. newspapers who were working to leverage a variety of different types of bots—automated software programs used for all sorts of tasks online—for their own purposes: to write stories, scour databases, and post updates to social media. All of these conversations—and the time spent in the field with campaigns, journalists, and digital constituents (U.S. citizens active in online politics)—had shown me the

innovation, chaos, and ruthlessness at the heart of digital political communication during the 2016 U.S. presidential election.

At the event in Silicon Valley a well-known data scientist from one of the world's largest software firms echoed the idea that the campaign had been particularly chaotic and innovative. He said that politics online was the "wild west"—that it was an anything goes space devoid of regulation. This phrase, "wild west", gave me a massive sense of déjà vu. Looking over my field notes later that night I realized that I had heard that exact sentiment—that the internet was still the "wild west" when it came to U.S. politics—from two other key informants during the contest. One was a social media marketing expert who had worked as a consultant on the Trump campaign (interview, February 2, 2016). The other was a reporter at one of the nation's most circulated newspapers (interview, April 18, 2016).

By election day 2016 a great deal had changed since Obama's noted digital campaigns in 2012 and 2008. Specifically, there was little to no conversation about the role of either political bots or disinformation online during these contests. More generally, the norms of political communication changed in 2016 because of the tactics of Trump and his digital director, Brad Parscale. Parscale was blunt about the importance of social media in 2016, consistently noting that platforms like Facebook and Twitter were crucial in helping Trump's campaign to massively scale support amongst the U.S. populace and in raising a great deal of money (Ellyatt, 2017).

B. Bots, Computational Propaganda, and the 2016 U.S. Election

This study focuses upon the communicative role of bots and computational propaganda during the 2016 U.S. presidential campaign. In terms of how it relates to the sending and receiving of information, Computational propaganda "describes the use of algorithms, automation, and human curation to purposefully manage and distribute misleading information over social media networks" (Woolley & Howard, 2017). Specifically, this study seeks to build understandings of the role of bots and computational propaganda via research alongside and amongst to three actor groups who built and used bots during the 2016 contest: campaigns/candidate, journalists, and digital constituents.

More than nine months of field work and interviews with over 40 experts from these groups provides insight into who built and used bots and computational propaganda, why they did it, and how they did it. This ethnographic information is crucial in generating understandings not just about the general use of these communicative technologies during the 2016 election, but also about how these tools are related to the people who build and use them.

Two broad questions are at the center of this research project. First, who made and used political bots and computational propaganda during the 2016 U.S. presidential elections? Second, what impact did builders believe these tools and strategies had upon during the contest?

This research is informed by theories, and speaks to literature, from both political communication and science, technology and society studies. By focusing on two

traditional actor groups from political communication (campaigns/candidates and journalists) I work to address the ways known actors within this sub-discipline of communication engage in the use of novel and, at times, manipulative technologies in their work surrounding elections. This research has also spurred the creation of a new actor group, digital constituents. Never before have digitally engaged citizens—and not only powerful political actors—used bots and computational propaganda in attempts to be heard during an election.

By focusing upon the communicative use of a particular technology, bots, and a particular informational strategy, computational propaganda, I speak to problems at the intersection of human-machine interaction. Specifically, I discuss how automation—and the bot as a proxy of the builder—allow those engaging in political communication to amplify their own content while simultaneously suppressing that of their opposition. I also discuss, however, how these tools are used by journalists and digital constituents in attempts to support or bolster democratic communication. I account for the ways in which these groups use bots and, more generally, automation on social media as a prosthesis for social organizing, reporting, and data parsing.

C. Manufacturing Consensus, Information Radiation, and the Bot Proxy

Three working theoretical concepts have grown from this research. Each coincides with research amongst one of the three actor groups. Data about the usage of bots and computational propaganda amongst candidates and their campaigns revealed a novel communication tactic for scaling political support online: manufacturing consensus.

Conversations with journalists consistantly surfaced the ways in which bots act as useful radiators of information for reporters who are otherwise engaged in story development. Finally, fieldwork amongst digital constituents—but also the ways in which campaigns and journalists used and deployed social media boys—showed how bots are related to, but not the same as, their creators. The bot is a proxy for the creator. The person, or people, who built a bot are responsible—to a degree—for what this automated software does. However, when a bot is deployed online it exists in ecosystems where inputs and outputs from other sources can result in unforeseen, and unintended, actions and consequences.

D. Political Communication, Technology, and Society

Any study of media and politics should "foreground complexity, interdependence, and transition" (Chadwick, 2013, p. 4). Evolutions in the technical and social spheres have led to the construction of a "hybrid" media system that has altered the dynamic of political communication. Socio-technical innovation has led to ever changing political affordances of this multi-media landscape: the uses of the internet, networks, algorithms, bots, online data caches, and more. The rise of hybridized technology and "networked society" have not only affected the way political conversations occur, they have also altered the ways campaigns are organized, elections function, and power is exerted (Benkler, 2006). New political organizations have been birthed, political systems have changed, and politicians have risen and fallen. Some aspects of political communication, however, remain constant.

This research project exists at the confluence of ever evolving hybrid technology and a society wherein communication is largely undertaken and defined via digital networks. By studying the people who build and use bots and computational propaganda, I work to answer crucial questions about how automated tools, social media platforms, and online tactics ranging disinformation to political harassment—often levied via these very tools and platforms—result in both communicative innovation and political normalization of new technologies. As Kaprf (2012) points out, "even the most radical changes to communications systems must be channeled through [standing political] structural constraints in order to impact traditional political outcomes" (p. 10). Many old guard members of the political elite remain the same worldwide, and newly powerful individuals and groups have ascended. These actors have adjusted, or made use of, the altered state of political communication as it exists online. In some ways, digital democracy has not played out as cyber-optimists had hoped. The elite on the internet are still elite and thus "online speech follows winner-take-all patterns" (Hindman, 2008).

Nowhere is this clearer than in the usage of political bots and computational propaganda during the 2016 U.S. presidential election. Bots played a crucial role in amplifying fringe views, but they also disproportionately amplified the ideas and positions of Donald Trump (Kollanyi et al., 2016). But my research reveals complexity within this communicative normalization. Digital constituents—regular people who learned how to use bots and computational propaganda to further their views and the views of the candidates they supported (or to attack those of the candidate they opposed)—used new communication tools in novel ways during the 2016 contest.

In some ways, this groups' use of bots simply worked to bolster the platforms of the powerful. At times during the election, regular people were tricked by fringe partisans into using bots and other tools to bolster trending topics that benefitted candidates and campaigns (Schreckinger, 2016). They were, in other words, tools of astroturf politics. In conversations with other digital constituents, and even with many journalists, it was clear to me that the very same communication tools were used in ways antithetical to political normalization.

This said, and despite the rise of a multi-platform and multi-mediated experience, the traditional concerns of political communication still have relevance (Chadwick, 2013). Both older and newer media matter, but distinctions between them, and between traditional and non-traditional political actors, are blurring (p. 4). The classical political communication study of the interactions between political actors, media, and publics should provide a foundation for beginning to study both modified and new units of analysis, according to Chadwick. Contemporary political communication is in a transitional space, but perhaps this is simply a facet of a field that makes its research focus the ever-evolving landscape of the problems surrounding interaction and power. The units of analysis for this subject are similarly in a state of revision and adaptation.

Communication between political actors, media, and citizenry must still be interrogated, but it should be mapped onto new and emergent ways of interacting. Studying countries, individual political actors, and clusters of citizenry and politicians is still a worthwhile

endeavor but such an exercise becomes more nuanced when considered alongside hybrid techno-social formations. Scholars of political communication must focus their study upon new units of analysis in order to better comprehend the current state of affairs. This research has revealed, time and time again, that it is integral for the discipline to investigate the political role of software and algorithms and of social networks.

E. The Politics, and Relevance, of Bots

The results of this study strongly support research by scholars of communication which suggests software algorithms are inherently political. As Gillespie (2011) points out, algorithms frame information in a particular way. Bots, by extension and much more obviously, do the same. The ways in which algorithmic framing of content on social media occurs is often complex and sometimes unknowable to even the coder who builds particular software. Bots can also, by way of anonymity, be difficult to detect and their effect upon mechanisms including trending algorithms can be opaque. Despite this, my informants in this research made it clear that they believed algorithms and software code can work to enforce a particular politics. Algorithms, software and bots are now communications technology, like broadcast television or online newspapers, and are "caught up in the ways we ratify knowledge of civic life" (Gillespie, 2014, p. 167-168).

With this in mind, this research prioritizes the interrogation of not only the people who build bots and used them in and around communication about the 2016 U.S. presidential election, but also of political bots themselves. Throughout this research project I followed and catalogued the activities of political bots identified by informants, other researchers,

and myself. One might call this ethnography of technological actors. I myself think of it as ethnography of techno-human culture—a hybrid culture that is informed by actors driving social change that are not always, or necessarily, human.

The unit of analysis for bots, software or algorithms can be chunks of script. It can also be, as is the case for this study, social bot accounts that exist on the front-end infrastructure of social networks—on Twitter, Facebook or Instagram, for instance. These accounts mimic real human accounts and can be tracked, followed, and researched. Their communication can be catalogued and anlayzed. But this work is generally much more useful, and communicative intentions are much clearer, when the study of particular bot accounts is paired with field work with those who have built said accounts. Whenever possible, I have done this in this research.

This study establishes political bots, and political botnets (collections of political bots), as units of analysis amongst scholars of communication. I argue that politicized social bots are among the most important developments in the technologically oriented study of political communication and suggest that social science scholars from all disciplines should work to understand the ways these coded political actors are society worldwide.

The classical study of communication and politics must work to understand the ways new technologies—including bots but also other emergent digital tools-- are used by emergent political actors. As Chadwick (2013) argues, "hybrid thinking rejects simple dichotomies, nudging us away from 'either/or' patterns of thought and towards 'not only/but also'

patterns of thought" (p. 4). This inclusive notion of political communication can only be useful to those attempting to understand today's complex communication landscape—both in the U.S. and abroad.

F. Summary of Chapters

The following pages are separated into five distinct chapters (beyond this introduction). First, I outline the methodology for this study. I describe my usage of field research methods to develop understandings about the use of bots and computational propaganda during the 2016 U.S. election and detail empirical specifics. Second, I discuss findings and analysis from research into the first actor group in this study: candidates and campaigns. Third, I explore how journalists used bots and computational propaganda—and were targets of these tools and strategies—during the campaign. Fourth, I describe and analyze the usage of bots and computational propaganda by digital constituents. Finally, I provide a conclusion that details methodological constraints of this study, paths for future research in this area, and frontiers of bot usage in and around political communication.

In chapter one, on methodology, I foreground the groups that I studied during this project and why. These groups are the makers of political bots, in-house engineers and industry data and computer scientists. The first group is spread across all three actor groups covered in chapters three, four and five—or candidates/campaigns, journalists, and digital constituents. The second group and third group largely relate the latter two actors. Though some campaigns are now bringing data experts, and even computer scientists, in-

house, most work with these professionals in on ad hoc contract or sub-contract projects. In the methods section I also explain the reasons for using networked ethnography, or fieldwork occurring across multiple sites necessitated by the transient nature of political campaigns and the communication associated with them (Howard, 2002). Finally, I discuss the nuance of using field methods to study both human and technological actors and discuss the literature and practices that informed the portion of this research housed within the growing field of ethnography of information.

In the third chapter I begin with an examination of the process of studying the use of bots and computational propaganda by candidates and their campaigns during the 2016 U.S. presidential election. I explore root concepts from communication and media studies, namely agenda setting, and explain how they are challenged by the rise of bot usage and computational propaganda in political communication. Through data from interviews and fieldwork I establish the concept of manufacturing consensus, a play on Herman and Chomsky's (2002) theory of "manufacturing consent".

While their work builds out a propaganda model that focuses upon the ways in which mass media are complicit, and even engaged, in control of public opinion, my research—and the notion of manufacturing consensus—relates to the ways bots are used to amplify content around candidates, campaigns, and political ideas in order to give them the illusion of support or consensus amongst the broader public. Crucially, I argue that this usage of autonomous political "supporters" on social media often affects the way the media frame conversations about candidate viability and, therefore, public opinion.

Chapter four is has a dual focus upon how journalists and media makers use bots in as tools for reporting—as what I deem information radiators. Importantly, my fieldwork with digital constants reveals that the use of bots by journalists is at times, by some informants, considered to be a version of computational propaganda. This grows out of digital constituents' perceptions about the subjective and partisan nature of certain elements of reporting. This chapter also explores the ways in which journalists in the U.S. were targets of computation propaganda. My research suggests that many of these attacks were undertaken in order to silence journalists by either spamming them over social media or scaring them into not reporting upon particular stories. I catalogue one particular instance of chilling effects due to computational propaganda against reporters during the 2016 contest. In this circumstance, bot makers and political trolls using a collective hashtag, #thelist, targeted Jewish journalists in a semi-successful attempt to silence reporting about alt-right anti-Semitism.

Digital constituents, the actor group discussed in chapter five, are explained as an emergent group that arose in this study through grounded observations during fieldwork. Data from this research suggests that this group is, in some ways, producing a democratization of not just political bot usage, but also computational propaganda. Many people I interviewed were regular citizens who used bots and social media in attempts to prioritize, and amplify, political content they saw as being most important. Discussions with members of this actor group were particularly illuminating when it came to exploring the nature of the relationship between bot builders and their bots. In this

chapter I explore the idea of the bot as a proxy for its creator—an anthropomorphized tool that can drive sociality in ways both related to, and unforeseen by, bot builders.

Finally, chapter six concludes this study by discussing methodological constraints of this research, areas for future research, and frontiers of political bot usage and computational propaganda uncovered since the original research during the 2016 election. The ways in which computational propaganda is different from traditional theories of propaganda—relating to media from books to broadcast—are unpacked in this chapter. Namely, I argue that automation and anonymity make computational propaganda a different—and more potent—version of an old political communication strategy. I also discuss elements of communication theory integral in this study and explore the ways in which different epistemological stances may be useful in producing future work on related subjects. I end this chapter with a discussion about the ways this research revealed bot builder's desire to confuse rather than control. Ultimately, I suggest that confusion is, in many ways, another path to control.

Chapter 2

Methodology:

Bots, Computational Propaganda and the 2016 Campaign

A. Studying the Builders of Technology, and Their Tools

During my work on the 2016 U.S. presidential campaign, and throughout my broader research on bots and computational propaganda, there has been a common misconception. Namely, that my work is focused upon analyzing large sets of data from Facebook or Twitter rather than that my contribution to these subjects is based on data and analysis from ethnographic methods. This has led to interesting occurrences in my professional life, but also during fieldwork itself.

One evening early on in my data collection I was speaking to an employee from the Clinton campaign in Detroit (interview, March 8, 2016). I was there for the Republican and Democratic primaries in Michigan, working to gather preliminary data on how candidates and campaigns might be making use of bots or other elements of computational propaganda. While describing my interest on the usage of bots during the election to the campaign worker I noticed his eyes begin to focus off in the distance beyond my head, a sign he was less engaged than I might have hoped. After I finished, he told me that he could nothelp me with research on computer science. No, I told him, I am a social scientist—I want to know what you know about this topic. That got him talking.

The findings of this project are informed by a variety of different sorts of fieldwork: networked, information-oriented, and, of course, elements of traditional ethnographic study. My goal for this project is to study the people who build, deploy, interact with, and track bots and computational propaganda in order to understand what they think about this topic. Why do they do what they do? What do they do? How do they do it? Who are they doing it for? These are all central questions to this work. I see this project as complementary to big data analyses of information from social media. Whereas that work illuminates details about, say, the quantity of political bots, their networks and their communication, this project focused on the quality of political bot—and computational propaganda—usage. Here I work to qualitatively interrogate the role of this phenomenon within political communication, and in U.S. society, during the 2016 election.

B. Introduction

Over the course of the 2016 U.S. presidential campaign and election I studied the ways that three actor groups used digital media, bots, and automation using a combination of field research methods. These groups were political parties and their campaigns, journalists, and digital constituents. From the beginning of February 2016 through to the weeks after the election that November—a total of approximately ten months—I used a combination of observation, interview, participation, and process tracing to build understandings of the campaigns and their digital maneuvers. The goal was to create a diagnostic, humanized, view of the way that people affiliated with the parties made use of bots and computational propaganda.

This project used field methods to gather knowledge of how bot designers operate professionally, both in terms of working as an innovative network of engineers and as a professional network competing for clients in a market for computing services. It is grounded within study of these networks during the 2016 Presidential Elections. I interviewed and observed three broadly categorized groups:

(i) ***Makers of Political Bots***. Political parties/campaigns, journalists, and digital constituents use bots in unique ways in their various attempts to spread information. This is true whether they are explicitly working to manipulate public opinion or if they are using automated software to search for and parse information in a broad sense. Some of these groups employ in-house coders to make their own bots. Contracted individuals and firms are used, however, in many cases to construct, deploy, and manage bots. In fact, some of the engineers behind political bots actually work for—and on behalf of, as sub-contractors—digital political consulting firms. I interviewed and conducted participant observation with makers from these groups.

(ii) ***In-House Engineers***. All of the social media firms with high profile social media services employ computer engineers to detect bots. US political campaigns, particularly big presidential ones, now have teams oriented towards technology and communication. Teams from both firms and campaigns worked to track bots. The process usually involves some algorithmic identification of problematic accounts that publish too quickly, but with the growing sophistication of bots human confirmation is often needed before accounts are sanctioned or deleted. I spoke to several individuals working on such teams for this research.

(iii) ***Industry Data Experts and Computer Scientists***. There are a few third-party organizations dedicated to identifying and researching bots. A few work for online services such as Status People or research projects like the Truthy initiative at the University of Indiana. Some work at major publications in cities such as New York, San Francisco, and Los Angeles. There are also small research teams at Microsoft Research, UC Berkeley, and NYU that track bots. I built ties with contacts that track and use bots for data oriented journalism and political commentary.

For this work, I define bots as automated software programs used to online do tasks that a human would otherwise—without technological help—have to do manually. I use the

following definition for political bot: “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” (Woolley and Howard, 2016, p. 4885). I also define computational propaganda using the definition from Woolley and Howard: “the assemblage of social media platforms, autonomous agents, and big data tasked with the manipulation of public opinion” (p. 4886).

To define social media, through which most social bot action occurs, I make use of Howard and Park’s (2012) three-part definition. 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 form of personal messages, news, ideas, that becomes cultural products; and (c) the people, organizations, and industries that produce and consume both the tools and the content” (p. 359).

This research is concerned with the makers of media, here political bot producers, because the creative practices of makers—and the social practices of individuals writ large—are constitutive of the practices of large-scale institutionalism and organization (Neff, 2012). Howard (2011) summarizes the arguments of Castells to this end in saying that networked media research “must involve studying large organizations that build and manage media infrastructure (here-social platforms), the individuals who produce and consume content over media (here-botmakers), and the content that is produced and consumed over media” (p. 2).

Networked ethnography is useful here because it works to gather data on mobile actors and the myriad 'trace' data associated with their communication through all sorts of interconnected media (Howard, 2002; Geiger & Ribes, 2011). Howard's study of multi-site political campaigns, and constantly travelling campaign staff, highlights the role of understanding social interaction via networked means. The work of (Burrell, 2009) takes this argument further to suggest that fieldsites occupy virtual, physical and imagined spaces.

Bot makers and trackers, whether from political campaigns, news outlets, nonprofit organizations, or super PACs, must be studied as a networked actors, from multiple and often fluid research sites, for several reasons: a) they are moving alongside elections, following candidates and writing bots to speak to the latest issues or crises, b) their primary means of communication is digital, thus necessitating both online and offline interaction, and c) they are strategically dynamic, relying upon networked structure as a means of creating an air of legitimacy, of transmitting both selectively visible and potentially viral content, and of manipulating relationships. The one to one, and even one to many, communication models are far from adequate in capturing the networked communication of these actors.

C. Research Questions

This project seeks to develop understandings of this emerging communication phenomenon via interviews and participant observation of the people who make and deploy bots during the 2016 US presidential election. The different ways these actors use

bots, and intended outcomes of this practice, is—therefore—central to this work. It is informed by broad research questions: 1) who makes bots and why? 2) What is the intended impact of bots on public life? In this study, focused upon an election in the United States, these questions become: 1) who makes political bots in the US and why? 2) What impact do bots have on political communication in the U.S.?

Several more nuanced sub-questions are at the heart of this research on political bots and the 2016 U.S. election. How might this technology be constructed and deployed with the intention affecting public opinion, or behaviors like voting and civic engagement? How are bots used by or against other democratic institutions, particularly the free press and non-governmental organizations, to generate or influence content and communication? More generally, I am concerned with the role of bots as drivers of social change external to the intentions—or direct actions—of their builders. How are bots challenging traditional notions of agency in the field of science and technology studies and traditional concepts of actors in the political communication?

Because of this unique perspective on agency, it is crucial that this project ask questions about novel methods for studying bots and their builders: In what ways does studying political bots and semi-automated entities contribute to, and extend, existing ethnographic methods? What do internet-oriented companies, particularly those on social networking platforms like Twitter and Facebook, do to track and curate political content generated by bots and how can this information be leveraged in research?

D. Studying Bot Builders, Bots, and the 2016 Campaign

My goal for this project was, beyond understanding widespread bot usage during the 2016 campaign, to build an understanding of other novel digital campaign methods for communicating information. I paid particular interest to communication methods that made use of both computational tools (automation, algorithmic design) and an effort, often subtle, to manipulate public opinion. I consider these efforts part of the broader spread of computational propaganda. Tools for dissemination and obfuscation like political bots where of central interest to my research, but so where the strategies that these tools helped realize: the spread of false news reports, "shitposts"—or highly negative memes, and attacks upon journalists.

In order to understand where these tactics and tools originated, I spent time in several states during primaries, attended campaign gatherings, sat in on digital strategy workshops, and traveled back and forth to the party nominees' home turf events in New York City. I wanted to develop a sense of the culture of the campaigns, so I needed to do some "deep hanging out." Participant observation formed a portion of this work. In order to gain access beyond the hordes of volunteers, and to get in touch with those in the know about digital strategy and with the ability to make assertions about party and campaign strategy, I had to spend time meeting people and learning the structure of the campaign apparatus.

I volunteered and interacted with several of the campaigns. I used applications like MiniVan to canvass for Bernie--knocking on doors in NYC Chinatown and the Bowery. I

made calls and sent texts in Detroit for Clinton and even hung out with her campaign employees at the preemptive Michigan primary "victory" party that turned out to be a shocking precursor for the later electoral loss in that state. I corresponded with people working for the Trump campaign, went to campaign headquarters—and was turned away—twice, and talked to crowds of red-cap-wearing supporters at various events.

During the New York mixed party primary, I attended a very sparse meetup at a Chelsea tech store organized by a relatively unknown digital firm working for Ted Cruz. The company's chief data scientist and director of sales outlined, in deep detail, the firm's work in "behavioral analytics". The firm would turn up later in the campaign, and in many sensational media stories, when it began deploying its alleged "psychographic" digital tactics for the Trump campaign. It was named Cambridge Analytica.

In order to stay up to date with events I signed up for every mailer from the campaigns, I followed the candidates on social media and religiously scoured their messages and their metadata. I became very familiar with the regular re-tweeters and likers—especially those that showed signs of automation. I took screen shots of public content, made notations about what they showed. These, and other field notes, were stored using Zotero and Microsoft Excel. It is worth noting that political bots are often short-lived. They either fulfill their task, and are then taken down by deployers to avoid a trail, or are deleted by social media platforms because they violate terms of service—i.e. showing signs of spam or being used to harass other users. I have included several snapshots of these bots here in

order demonstrate particular tactics and types, but also to preserve now nonexistent automated accounts.

I followed important campaign events and important moments online when I could not attend in person. I used news reports, community documents, and archived social media material to understand such events. I also gathered second hand accounts, via one on one interview, from experts who had been in attendance or who had worked with or for the campaigns. Contradictions in stories of how events played out, or in how automation or other social media tools were used, regularly occurred. This was a highly-contested, and strategically ruthless, campaign. Parties and campaigns, and even factions who worked within them, regularly disagreed about how things happened—about what truth looked like. In order to get to the bottom of discrepancies in accounts, I worked to complement them with online narratives from multiple sources. I captured these, and all other sources, using Zotero.

Figure 1: @DyanNations Bot

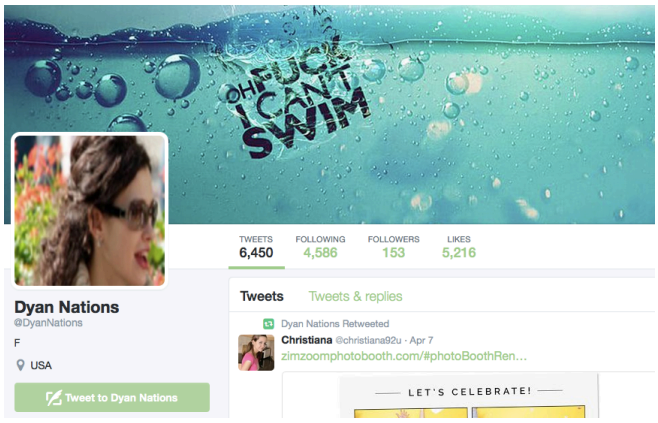

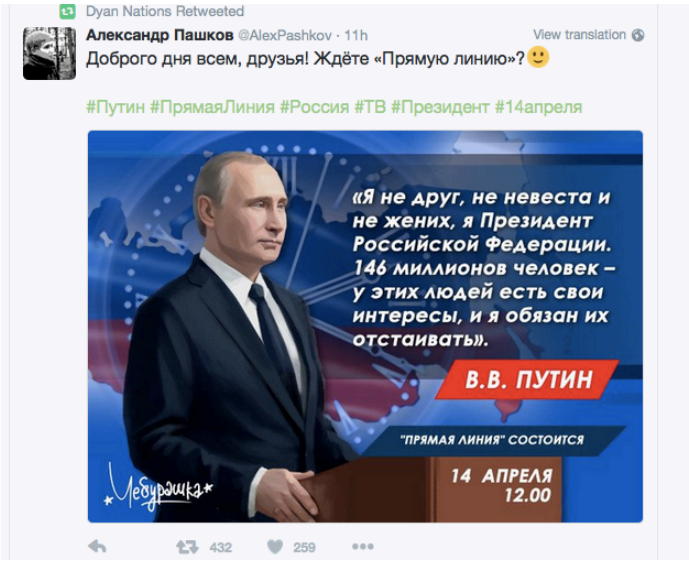

Source: Author's screen capture, 6/15/17.

Note: Two screen captures from @DyanNations, one of several hundred accounts conservative strategist Patrick Ruffini alleged was used to attack Ted Cruz on behalf of Donald Trump. In addition to tweeting pro-Trump messages, the account regularly sent out Russian memes and ads for fake followers.

E. Field Methods and Technologically-Oriented Communities

There is a short, but rich, history of using ethnographic methods, including interview, to study technology-oriented communities. The proliferation of digital qualitative methods has grown out of the work done by the social scientists and humanistic scholars throughout academic history. Scholars have struggled with both the proper terminology for studying culture online, vacillating between methods that are digital, virtual, net-based, 'netnographic', and data-centric, but also with how to represent culture online (Bowler, 2010). Those interested in internet methods early on were particularly concerned with the separation between the 'real' world offline and the artificial one online. This separation itself has since been criticized as artificial, especially in light of the ever increasing 'wired' nature of society.

Markham and Baym (2009) strike down the notion of online/offline separation by arguing that sectioning off certain aspects of social life as real or digital is a colloquial endeavor, one that has no place in empirical research. They argue, moreover, that "qualitative research requires a tolerance for chaos, ambiguity and inductive thinking" (p. ix). By being open to the potential for multiple, and concurrent, meanings—to subjectivity—qualitative researchers are able to work towards contextual understandings of new phenomena as well as new ways of thinking about existing phenomena. This perspective is crucial when theorizing about the social capabilities of technology.

Schools of thought from actor network theory to activity theory seek to answer the complex questions associated with the ability of technology to affect change. Social bots are an obvious object of analysis within the category of technology with a social function. The concept of the bot as a proxy allows the simple polarization of the technical and the social to be complicated. My conversations with bot builders made it impossible to simply classify bots as the technical outcomes of social creators. Bot builders regularly negotiated their relationships to the bots they built. Their expressions of reflexive identification and moments of comparison between themselves and these automated social programs illuminated the proxy nature of bots.

Interviews, and other ethnographic methods, allow Communication researchers operating within the subfield of Science and Technology Studies to get to the heart of intentionality. They ask technology makers why they built what they did, how they built it, and what they believe it does. This strain of research is in contrast to user-based

studies, which seek user reactions and interpretations of media. This is not to say that the two cannot be complementary. Here, however, we are concerned with the relationship between builders and their bots, so we have conducted over 40 interviews with the bot building community. In this paper we focus spoke mostly to builders who make Twitter bots, but this research has implications for other social platforms where bots are being used: Facebook, Reddit, Instagram, Snapchat, Kik, Periscope, Firechat, Wechat, Telegram, Slack, Weibo and many others.

My interviewees come from a diverse range of professional, demographic, and cultural backgrounds. Some of them work as professional software developers; others are digital artists, educators, lawyers, journalists, political consultants, entrepreneurs, or writers. Many work day jobs that have no connection to bot development or the digital sphere but engage in bot building as a hobby or activist practice. I spoke to men, women, and those that self-identified as non-binary. Some were in their late teens or early twenties, others were in their 30s, 40s, and 50s. Politically, interviewees spanned the range from conservative to liberal. Some have worked for political campaigns, others consider themselves anarchists. One thing brings this otherwise heterogeneous group together, they all build bots. The ideas in this paper, therefore, come from technology makers.

The interview process for this project includes four primary steps: 1) establishing contact with potential interviewee, providing relevant information, and asking for consent, 2) conducting the interview, 3) taking notes, and 4) writing a thematic memo for each interview. I audio record all interviews, unless asked not to by interviewees. I fill out an

interviewee information statement for every interview. This sheet includes demographic information, but also general questions about the person's level of coding experience and professional background. I give all informants pseudonyms.

This was a considered choice made in the beginning stages of my project and continues for two main reasons: 1) to protect informants' identity, especially when confidential or politically sensitive information is being shared and 2) to allow for open conversation, so that thoughts and feelings can be shared without hesitation. In order to protect this data I abide by ethics guidelines established by both the University of Washington and the University of Oxford, but also those of funders including the US National Science Foundation and European Research Council.

F. Studying Bots as Drivers of Sociality

Hine (2005) makes an early effort to ask, and work to posit answers to, pressing questions faced by scholars doing work at the intersection of technology and the social world. What does internet mediation do to human interaction? Do old methods of qualitative inquiry have use? If so, how can they be amended to better facilitate online research? Are new methods needed? Ultimately, Hine suggests that the stress and excitement surrounding the study of interaction via a new medium be grounded in the idea that all qualitative method is contextual to the site of study and must be adapted *in situ*. She argues that methods, far from being apolitical, are encoded with meaning by both researchers and publics.

This representational move makes important distinctions about the way qualitative scholars choose to speak about sociality, whether technical or human—whether social bot or machine learning algorithm. Hine contends that qualitative social researchers should be open to the ways technology might alter the social science, both epistemologically and methodologically (p. 9). This move creates a path for reflexive theorizing about how humans affect technology but also about how technology affects humans.

Markham and Baym (2008) argue that "qualitative research requires a tolerance for chaos, ambiguity and inductive thinking" (p. ix). By being flexible qualitative researchers are able to work towards contextual understandings of new phenomena or new ways of thinking about existing phenomena, such as is the case with theorizing the social capabilities of technology. Bots fit nicely into the category of social technology. It is Markum and Baym that argue that research requires the ongoing balance of different dialectical tensions and that researchers must work to articulate what is gained and lost by particular arguments (p. xix). With this in mind, the polarization of the technical and the social can be complicated by way of the process of reflexive identification and comparison.

The argument that "technology is not one more iota mediated by the digital" is central to the work of Horst & Miller (2013) and suggests that the notion of greater pre-digital authenticity is a fallacy (pp. 3-4). As the scholars who pioneered the University of Chicago method of qualitative analysis suggested, that which is perceived as real has real consequences. For Horst and Miller that interaction which occurs via technology is still

tangible, and fundamentally social, pointing towards the idea that the social permeates all facets of the world as humans see it, mediated or no. This ontological point is important in theorizing the role machines play in knowledge generation but also in directing qualitative methods for understanding technical society towards greater reflexivity. While Horst and Miller advocate for a reconsidered multi-methods approach that well suits the networked nature of bot making community, the messages they produce and the secondary data that covers their actions.

The concept of trace ethnography, presented by Geiger & Ribes (2011), is based in this idea that ethnographic researchers ought make use of the multitude of data online, but also that we must interrogate the unique ways that information technology systems drive both communication and, by proxy, sociality. Geiger and Ribes argue that this tactic "combines the richness of participant observation with massive amount of data in logs in order to reconstruct patterns and practices of users in sociotechnical systems" (p. 1). Ethnographers of technology-oriented communities must find ways to work with trace social data in tandem to traditional fieldwork practices in order to comprehend the heterogeneous nature of the networked environment (p.1). Beyond this, they must think about how particular eco-systems online cultivate particular types of communication while prohibiting others.

Coleman's (2014) multi-year participant-observation amongst the activist-hacking group Anonymous underscores this relationship between human individuals/groups that form the larger organization and the digital technology that facilitates their practices. For

Coleman, Anonymous would not exist without technological platforms like 4chan and indeed a wide array of other computational tools. She suggests however, that technology, and especially the infrastructure of internet security, has been significantly changed by Anonymous. This relationship is indicative of the larger goal of ethnography of information in that it reveals the interactive social relationship between both machines and humans.

Though ethnographers of technological communities often work to understand emergent phenomena, the social role of technology, they must still obviously work to capture interaction. Qualitative research has been built to understand human society and culture, and the communication therein, but smart tools are only understood as such because they have emerged via interrelations with humankind (Latour & Venn, 2002). With this in mind, I make use of classic and emergent qualitative methods texts in order to give intellectual breadth to the technological depth of this study.

This ethnography of information is a challenging practice because it requires the researcher to prioritize the study of non-human actors as contributing to larger patterns in communication and society. Qualitative interrogation of bot accounts on social media is useful in that allows for deep study of how particular automated accounts might affect their ecosystem is unusual ways. When the researcher knows the person that build the bot in question, then they can compare what the bot does with what the builder intended it to do. These things are often at odds. But it isn't always, or even often, the case that we can know the human behind the bot. Beyond this, bots do interesting things on social media

outside of their relation to their builders or even to people. They communicate with other bots, they communicate with trending algorithms—especially when they congregate and communicate in large numbers about a particular issue.

Social media bots, and political bots, are often built to look like people. But they aren't people. They are encoded with values by their builders but they also operate in ecosystems that prioritize the inputs of the bot's creators alongside those of other users in the system. This results in unexpected outputs from bot accounts—the case of Microsoft's Tay being one among many. Quantitative research is useful in studying how social media bots—or highly automated accounts on a site like Twitter—are used to communicate broadly speaking during a particular event or block of time. This work can deduce numbers of automated accounts and how many were speaking about which thing. Qualitative research, and ethnography of information, is more useful at deep study of how non-human actors can drive sociality.

New methods are needed to study how these "tools", automated or not, interact with humans and one another. It's not enough for scholars of communication, particularly those studying political communication, to only examine the spread of information between people. This is especially true in a world where digital communication and unseen automated processes dictate the exertion of power. It's true that well-resourced and well-educated human actors can more effectively leverage these systems in order to manipulate other people, especially during elections and security crises. It is also true, however, that systems constructed by large groups of people that are constituently

moving on to other work—Facebook's Newsfeed algorithm, for instance—are both unknowable by one person and do things outside the purview of any group of controlling engineers at any given time. When billions, millions, or thousands of people interact with an algorithm built to learn from its surroundings surprising things happen. Even bots that are built to simply message the same content repeatedly can have unexpected effects on communication and society when released into ecosystems of many other people and bots.

Chapter 3

Bots and Agenda Setting in the 2016 U.S. Election: Candidates and Campaigns on Social Media

"The point of modern propaganda isn't only to misinform or push an agenda. It is to exhaust your critical thinking, to annihilate truth."

-Gary Kasparov (as told to Sykes, 2017)

A. Introduction

On the evening of November 8, 2016 Manhattan was uncharacteristically quiet. At 10 pm I walked from an election return watch party at an office in the Flat Iron District to my hotel a few blocks north in Chelsea. Looking through the windows of shops, bars and restaurants of the 6th avenue, I saw hushed crowds of people with eyes to glued TV screens. I passed by a crowd of laughing men wearing Make America Great Again hats. Republican candidate Donald Trump had begun to shift towards the lead in the swing states of Florida, Ohio, Michigan, Wisconsin, and Pennsylvania among them. My contacts at the Javits Center, where Hillary Clinton was hosting her election night gathering, told me that the scene had changed markedly from the jubilation and confidence of just an hour before. People were beginning to leave in droves and several of the scheduled speakers had reportedly not taken the stage. The candidate herself was notably absent. Around 2:30 am the Associated Press called Wisconsin for Trump and,

with this, projected that he would win the election. It would, without a doubt, go down as one of the most unexpected presidential victories in U.S. history.

Most pundits, pollsters, and political scientists were wrong in their predictions of a strong Clinton victory. The day after the election GOP Strategist and "Never-Trumper" Mike Murphy wrote on Twitter "I've believed in data for 30 years and data died tonight. I could not have been more wrong about this election" (Cancryn, 2016). Months later statistician Nate Silver, whose website FiveThirtyOne had given Clinton a three-times greater chance to win then Trump, would disagree with the widespread sentiment that the election spelled the "death of data" in politics. In an interview with the Harvard Gazette Silver said:

> One thing I think is a myth is the notion of, 'Oh, polls got the election wrong.' I do not think that's true at all. The polls pointed toward a competitive race […] I think it more reflects blind spots in people's thinking. (Pazzanese, 2017)

Throughout the last months of the election, Trump and his campaign staff were dismissive—correctly—of media narratives and poll numbers that suggested they were far behind.

I argue that it was not as simple as data being flawed, polls being wrong, or media groupthink causing imagined consensus about a Clinton victory. In reality, the landscape—in terms of both data and socio-cultural perspectives on the election—was incredibly complicated. Assumptions about democratic prowess in the media tended to ignore the inner mechanics of the electoral college. Pollsters and journalists failed to

translate the ways in which a 30% chance of victory for Trump, and consistently close national polling averages, showed a statistically less then certain race.

The numbers on social media, particularly on Twitter, also showed a different election. Trump consistently received a larger market share of attention on the site than Clinton (Coyne, 2017). The interaction metrics he achieved on the site—follower numbers, retweets, mentions, and likes—also outpaced his Democratic rival (Y. Wang, Li, & Luo, 2016). Trump himself regularly pointed to Twitter popularity as factor in his win. In a March 2017 interview with Tucker Carlson Trump said, "I wouldn't be here if it was not for Twitter" (Fox News, 2017).

But if something was not quite right about the survey data that the public pollsters used to predict a Clinton win, the Twitter data was seriously flawed when it came to discerning candidate popularity. Though research by Wang et al. (2016) showed "Trumpists" on Twitter as being more engaged and more influential than "Clintonists," it relied upon face value reports of the follower numbers of those Trump followers to determine their influence. When Republicans, including Trump himself, referred Twitter as a marker of a public appeal not represented in the media they also tended to point to his follower numbers. In January 2017, CNN ran a story about the then President Elect reaching 20 million followers (Stelter, 2017). Reporter Brian Stelter wrote, "the president-elect brings up his social media stats on a regular basis, adding up the users who follow him on Twitter, Facebook and Instagram to say he reaches 46 million people through the three sites."

In the same story, however, Stelter (2017) made a crucial note, “The total [number of Trump Twitter followers], while impressive, is belied by the fact, like many social media users, Trump has some followers that are either robots (known as "bots") or dormant users who never see what he posts.” In fact, both bot detection tools and the researchers who use them suggested that Trump’s overall follower numbers on the social media were around 50 percent fake (Bessi & Ferrara, 2016a; Ott, 2017). Reports claimed that Clinton had a similarly high number of bot followers on Twitter (Johnson, 2016). Because swathes of bots follow political candidates and correspondingly drive up other metrics associated with political content (likes, hashtags, retweets, replies), using Twitter as a measure for political support is inherently flawed. But it does not mean that bots do not have influence upon perceptions, both in the media and publically, of whether or not a candidate has political cachet. On the contrary, media outlets regularly refer to twitter—and the numbers therein—when deciding on what story to report (Farhi, 2009; Vis, 2013).

In this chapter I discuss the ways that candidates, campaigns, and campaign associates used and interacted with bots and computational propaganda during the 2016 U.S. presidential election. Contrary to received wisdom, which suggests that the most invasive political bots, disinformation and computational propaganda came from the Russian government or groups such as the AltRight or AntiFa, I argue that the use of these tools and tactics has been—an remains—a strategy of mainstream political campaigns and their subcontractors in the United States.

Several questions come up when considering the role of bot builders, particularly those affiliated with political candidates, during the 2016 U.S. election. Did builders use bots to generate substantive fake online political traffic? What about subsequent offline support for the candidates? Who was buying and releasing the bot accounts that tweeted in support of Trump or Clinton? Who was building them? How did the campaigns view usage of bots? Did they make use of them?

In the previous chapter, I outline the ethnographic methodology used for this study. This chapter explores the research questions above by accessing over nine months of field research data, including interviews with campaign staff, digital political strategists, and—most importantly—bot builders. It argues that political bots—social media bots used explicitly for the manipulation of public opinion—are now crucial tools for campaign managers hoping to get an upper hand on the opposition by manufacturing political consensus. It also makes the point the political efficacy of these software tools is now understood, and used by, broad elements of the public in the United States and that there has, with this, been a democratization of computational propaganda in this country.

B. Agenda Setting, Digital Communication and the 2016 Campaign

Halfway through nine months of fieldwork on the 2016 U.S. presidential campaign trail I had an ephipany. I was talking to Cassidy, a digital strategist who did contract work for the Trump campaign, when he used the jargon of the discipline of Communication to explain the unlikely political ascendance of Donald Trump (interview, September 16,

2016). Cassidy brought up agenda setting, a theory that suggests that the more often something comes up in the media the more likely the public is to consider it important (McCombs, Shaw, & Weaver, 1997). Agenda setting is generally about power, specifically the power of the media to define the significance of information.

Cassidy said that the Trump campaign turned the concept upon its head. He said:

> "Trump's goal from the beginning of his candidacy has been to set the agenda of the media. His campaigns strategy is to keep things moving so fast, to talk so loudly—literally and metaphorically—that the media, and the people, can't keep up."

Here we learn that the campaign had a clear goal of attempting to guide and control what the media reported upon. Cassidy made it clear that Trump's campaign wanted to create scenarios wherein the media could not resist covering him and that this was a conscious tactic. Trump inverted, or perhaps twisted, the typical structure of agenda setting. Cassidy argued that the candidate's fast-paced rhetoric and willingness to speak "off the cuff" gave an impression of authenticity that demanded attention.

By defying expectations of what a presidential candidate should say and do, and doing so constantly, he worked to set the media's agenda which then, in turn, set the public's. According to a report from the Harvard Shorenstein Center on Media, Politics, and Public Policy:

> Overall, Trump received 15 percent more coverage than [Clinton] did. Trump also had more opportunities to define Clinton than she

> had to define him. When a candidate was seen in the news talking about Clinton, the voice was typically Trump's and not hers. Yet when the talk was about Trump, he was again more likely to be the voice behind the message. "Lock her up" and "make America great again" were heard more often in the news than "he's unqualified" and "stronger together". (Patterson, 2016)

Trump's bid to control the conversation, particularly in the mainstream media—but also on social media—paid off.

Cassidy and other digital strategists repeatedly stressed to me that candidates and campaigns are constantly working to stay up to date on a variety of evolving digital campaigning tools and a constantly changing communication landscape. Strategists associated with both the Republican and Democratic campaigns said that interactive advertisements, live-streamed video, memes, and personalized messaging all played a role in the spread of partisan content during the 2016 election. Nowhere is this clearer then across social media platforms and applications.

This is consistent with what other political communication scholars have found. According to Daniel Kreiss (2016), "while campaigns have long had to adapt to changing media environments, the pace, scale, and social consequences of change are qualitatively different in an era of rapid shifts in the application layer of the internet." The tacit goal of using digital tools, including infrastructural aspects of Twitter and Facebook, is to effect voter turnout. However, informants told me that elements of social media, including bots, were also used to achieve other, less conventional, goals: to sow confusion, to give a false

impression of online support, to attack and defame the opposition, and to spread illegitimate news reports.

One tool has risen to prominence amongst those used to achieve these latter means—to spread computational propaganda. That tool is the political bot. Previous research shows that political candidates and campaigns in both the United States and abroad have made use of these automated software devices in attempts to manipulate public opinion over social media (Ratkiewicz et al., 2011; Woolley, 2016). The 2016 U.S. election, however, was a watershed moment for the use of political bots and computational propaganda. Research from several sources suggests that political bot usage was at an all-time high during key moments of this particular election (Bessi & Ferrara, 2016; Howard et al., 2016; Ferrara et al., 2016). The Trump campaign, in particular, benefited from wide spread use of political bots. Supporters used these social automatons to give voters the impression that the campaign had massive online grassroots support early on, to plant ideas in the news cycle, and to effect trends on digital platforms.

C. Manufacturing Consensus and Democratizing Online Propaganda

In setting out the theory of agenda setting, Maxwell McCombs and Donald Shaw (1972) wrote specifically about the media's ability to teach voters via the information provided during the 1968 presidential campaign. The prescient line from Bernard Cohen succinctly explains the role of the media in prioritizing information: the press "may not be successful much of the time in telling people what to think, but it is stunning successful in telling its readers what to think about" (1963, p. 13). During a heavily contested

election, like the one in 1968 but also like the one in 2016, McCombs and Shaw argue that the power of the press in this capacity is significant. They argue there is a strong correlation between how often media outlets focus on particular campaign (or candidate) issues and whether voters think those issues is important (p. 181).

The question arises, however, who decides what the media reports upon? The traditional answer from the discipline of Communication is that gatekeepers, editors, editorial boards and the like, set the news agenda. The way this agenda setting process happened during the 2016 election was different, however. According to informants who did consulting work for the Trump campaign, the teams in the primary New York office and the digital office in San Antonio were well aware of their candidate's ability to capture and cast the news in his favor. The traditional campaign playbook was thrown out by this campaign in favor of one where all news on the campaign was good news, where a reliance of the short attention span of the media—and the ephemeral nature of stories in the 24-hour news cycle—meant that anything could be said and promptly forgotten. This caused Trump to draw extraordinary attention. It also caused him to gain an extraordinary amount of power in driving media coverage.

Digital communications strategy was a key part of the Trump campaign's victory in maintaining a consistently disproportionate share of media attention (when compared with the Clinton campaign's share, that is). Traditional media's willingness to cover Trump for free, to put him at the center of the presidential conversation, was one part of his success. The bigger portion, however, can be attributed to the way that he and his

supporters used social media. Twitter proved a crucial tool for Trump, a soapbox that allowed him to circulate content regardless of message.

Eastwood, an informant who has worked on several high-level Republican campaigns, was quick to point out that social media content was posted and then legitimized by constant coverage by major TV news channels, radio programs (Eastwood, 2016). Trump, he said, "won the election on social media." This win was also reaffirmed by hordes of political bots—automated social media accounts built to look like real users and used to artificially boost content about Trump with, according to Eastwood, the goal of manipulating public opinion.

In this chapter I argue that armies of political bots allowed campaigns, candidates, and supporters to achieve two key things during the 2016 election: 1) manufacture consensus and 2) democratize online propaganda. Social media bots manufacture consensus by artificially amplifying traffic around a political candidate or issue. This concept is unpacked in the following section of this chapter. To put it simply, however, manufacturing consensus occurs when armies of bots built to follow, retweet, or like a candidate's content in order to make them seem more legitimate, more widely supported, than they actually are. The goal of this it to galvanize political support where it might not previously. The illusion of online popularity for a candidate is created with the intention of spurring a band wagon effect. Trump made Twitter the primary communication channel in this election, and voters paid attention.

According to many of people I interviewed for this research project, including political bot makers and campaign personnel, the goals of bot driven tactics are manifold: to spur a bandwagon effect, to build fake social media trends by automatically spreading hashtags, even to suppress the opinions of the opposition. Some informants, including well known digital consultants, were dismissive of the effects of bots. They considered them "one tool among many" in digital campaign efforts defined by a common ethos of "throwing things against the wall and see what sticks" (interview, May 14, 2016)

However, they were also cagey and defensive about the political and moral legitimacy of using automated, and often fake, social media profiles to sway voters. In this defensiveness they overlooked, or failed to understand, the ways bots very genuinely effect political communication. Bots allow for the democratization of digital propaganda because they make it possible for one person or group to massively enhance their presence online. Open APIs, and laissez-faire approaches to automation on sites such as Twitter allow both campaigns and regular people to deploy their opinions large scale. As one informant, a political bot builder, put it: if one person operating one profile can automate their profile to tweet every minute, just think what one person running one thousand automated profiles can do (interview, February 2, 2016)

D. Manufacturing Consensus

The concept of manufacturing consensus is drawn not just from the ways bots and computational propaganda were used during the 2016 US election but also from a range of parallel uses in multiple countries dating back as early as 2007 (Robb, 2007; Gorwa,

2017). In the Americas, political actors in Mexico (Verkamp & Gupta, 2013), Ecuador (S. C. Woolley, 2015), Venezuela (Forelle et al., 2015), and Brazil (Arnaudo, 2017) have pioneered the deployment of bots in attempts to boost credibility through increased metrics: follows, likes, retweets, shares, comments, etc.. The ways social media platforms were used to manufacture consensus surrounding candidates and campaigns during the 2016 U.S. election, though, brought aspects of the phenomenon to the forefront of the global political psyche. Both the tools (bots), and the communication strategies they afforded (disinformation, algorithmic manipulation, boosting of social metrics), were catapulted to the forefront of reports from major new outlets from New York to Tokyo (Manjoo, 2017; Tanaka, 2017). Why, however, did this phenomenon become so much more widely reported upon during the Clinton-Trump election when it had been noted as a political communication strategy in previous global electoral contests?

First, Twitter played a more central role in candidate and campaign communication during the 2016 US Presidential election than it has in any political contest since the platform was launched in 2006. According to a Nov. 7, 2016 data report from Twitter, "people in the US sent one billion tweets about the election since the primary debates began in August of last year" (Coyne, 2017). Both party nominated candidates, but especially Donald Trump, used—and continue to use—the platform as a frontline communication channel to voters. Both before and during the election Trump pointed to his Twitter metrics as evidence that he had a backing not revealed in mainstream polls.

However, research reveals that communication of Trump related content was heavily bolstered by bots (Kollanyi et al., 2016; Bessi & Ferrara, 2016). The study by Kollanyi et al. found that, in the week preceding election day, bots messaged 5 times more for Trump then they did for Clinton. Moreover, an analyses of major primary candidates' Twitter followings by *FiveThirtyEight* found that Trump led the pack in fake followers (Bialik, 2016). A report the Twitter analysis tool Twitter Audit corroborated this assertion, finding that nearly half of Trump's Twitter following was fake (Bilton, 2016).

Figure 2: Trump Tweet About Social Media Following

Source: Author's screen capture, 7/24/17.

Note: Donald Trump consistently pointed to the size of his social media following as a sign of his wide-spread popularity. Several analyses from during the election, however, suggest that his Twitter following was around fifty percent fake—made up of highly automated or otherwise suspect accounts.

Data from Twitter's pre-election day campaign metrics report shows that Trump got more engagement—in the form of Tweets mentioning him—on the platform than his opponent (Coyne, 2017). News outlets consistently used Trump's Twitter popularity as a marker of his ability to control tenor of the election, citing metrics as evidence of engagement—despite the fact that such metrics are easily gamed through relatively simple automated code introduced through the API. This reification cycle, wherein bots boosted Trump-related content, Trump and his surrogates used the platform as a frontline communication channel to citizens, and media outlets reported on Trump numbers as sign of his popularity occurred consistently throughout the election.

The second reason key elements of computational propaganda and manufacturing consensus—especially political bots—rose to prominence in the global news zeitgeist was also borne out of the high-profile usage of Twitter and Facebook by the candidates, campaigns and their supporters. This very usage, coupled with previous knowledge of political bot campaigns during contested elections in other countries, led journalists and academics to closely scrutinize social media, not just as generally important communication channels but also as increasingly intertwined with official candidate statements, strategies, and fundraising efforts of the campaigns. Again, informants who worked on digital campaigning for Trump and other Republican candidates, made it clear that web-based marketing was a primary focus of the Trump campaigns' overall marketing and communication strategy. Indeed, it was the only media arena where the candidate outspent Clinton (Delany, 2016).

According to recent work by Kreiss (2017) on 2016 the campaigns' technology and media infrastructures, the Trump team treated the platforms and their representatives as consultants while Clinton's treated them as vendors. Elmira, a digital marketing specialist who worked for the Trump campaign, backed this assertion up (interview, November 16, 2016). She said that several of the major tech firms had an active presence in the Trump Digital office in San Antonio. Moreover, she alleged that the firms were engaged in active bidding against one another to get the campaign's business.

Conversations with members of Clinton's campaign team made it clear that the companies made similar efforts to engage with the democratic side too, but that they were much less receptive and more focused on building a traditional media campaign (interview, March 8, 2016). As Jessie, a Clinton campaign informant who worked on marketing put it, they were "working on Obama campaign model" when that model had either lost viability in the face of new technology and strategy or simply didn't fit the candidate at hand (interview, March 7, 2016)

The close relationship between Trump's digital team and employees of Google, Facebook, and Twitter meant that the team made use of each company's most recent marketing techniques. However, both Cassidy and Eastwood made it clear that less official efforts to massively boost the content of Trump and his surrogates, whether through automated likes, reposts or messages, undergirded the candidate and campaign's professional use of social media. Bot-driven communications allowed Trump to be perceived with a legitimacy and popularity that undoubtedly had effect throughout his

campaign. This is not to say that the social media companies condoned bot-driven activities. Rather, the other services platforms sold to the Republican candidate and his digital media team's receptivity to these services, combined with the less palatable bot-driven efforts at consensus building, came together to aid in the campaign communication efforts that won him the election.

Informants were emphatic that social media bots were an active part of previous US electoral campaigns on which they had worked. They were less forthcoming about whether either the digital arms of the Clinton & Trump campaigns actively sought out and bought bots to support their communication efforts. They implied that if this was done, it was done through hidden channels, but were not able to convincing evidence. Both Cassidy and Eastwood, but also Elmira—all of whom had inner working knowledge of the campaigns, suggested that methods for buying and deploying political bots were shadowed in the depths of campaign consulting efforts, rather than directly tied to the official campaigns and their staffs (interview, September 16, 2016; interview, May 4, 2016; interview, November 16, 2017). While the Federal Election Commission (FEC) tracks what campaigns spend, they do not track specifics related to funds that go to subcontractors. This layer of obfuscation means that all manner of communication tactics, from bot usage to efforts to propel false stories about the opposition, can be used by paid consultants at the tangents of the campaign apparatus without an obvious money trail.

Finally, computational propaganda, disinformation, and political bots—alongside other communication tactics for manufacturing consensus—rose to prominence in the 2016 election because of allegations of Russian State interference. Reports by multiple US Intelligence Agencies and foreign allies' counterparts said that there had been transnational manipulation of the US election by the Russian Government. Moreover, these accusations suggested that social media was a key channel for said manipulation. During the May 2017 Senate testimony of former Acting Attorney General Sally Yates' and former Director of National Intelligence James R. Clapper Jr. it was made clear that bots were a crucial part of the "Russian toolbox in the 2016 election" (2017). Senator Sheldon Whitehouse, D-R.I., said, "I went through the list [of tools used by the Russian's], it looked like propaganda, fake news, trolls, and bots. We can all agree from the IC (Intelligence Commission) report that those were in fact used in the 2016 election."

Manufacturing consensus occurs not only when bots boost social media metrics, but also when the media reinforces illusory notions of candidate popularity because of this same automated inflation of the numbers. It all, in turn, relates back to the fact that Trump and his campaign successfully took charge of agenda setting during the 2016 contest. As the *Business Insider* article on Trump's Twitter success put it:

> With each passing week, something [Trump] said, did, or promised to do dominated the news cycle. His words, however you felt about them, were consistently at the forefront of America's mind. And when you control the message for 17 months, you usually win. (Dunn, 2016)

It is worth noting that the way bots are built to engage with content extends from the passive, liking or retweeting, to the active, messaging or commenting on messages. This means that attempts to manufacturing consensus are not just driven by numbers, or vanity metrics, but also by active communication. While both actions mimic the actions of human users, the former are much easier to manipulate.

E. Media Agenda Setting from Political Bots

In order to understand the success of the Trump campaign's media strategy one need look no further than the early days of the campaign. In January 2016 Trump began gaining traction as a viable Republican candidate for the Presidency. In an opinion article for the New York Times written that same month, Peter Wehner, a senior fellow at the conservative Ethics and Public Policy Center and employee of three Republican Presidents, said he would never vote for Trump. He summed up the fears of the growing "Never Trump" movement taking hold within the Republican Party when he said, "no major presidential candidate has ever been quite as disdainful of knowledge, as indifferent to facts, as untroubled by his own benightedness" (Wehner, 2016).

Informants, including people that had done digital work for Republican presidential and senatorial candidates, saw Trump as a "loose cannon" willing to say and do anything (interview, May 4, 2016). They echoed Wehner's concerns about his lack of military or government experience. So did key members of the Republican establishment. Mitt Romney, the 2012 Republican presidential candidate, gave a speech in which he said: "Dishonesty is Donald Trump's hallmark […] He's not of the temperament of the kind of

stable, thoughtful, person we need as a leader. His imagination must not be married to real power"(Associated Press, 2016). Romney and his compatriots argued that it was only a matter of time before Trump did something so off-color that he would be drummed out of the race, but nothing seemed to be able to touch him. Media storms about Trump mimicking a disabled New York Times reporter, impugning Senator John McCain's war record, and harassing women only did not stick. Any one of these stories might have undone another candidate. Suddenly, however, Trump would say or do something else and a misstep would be forgotten in the next day's media cycle.

Experts from every quarter have since weighed in on what caused the Trump win. Communication scholars have suggested it has to do with the fact that, despite his disregard for traditional advertising and what his supporters have derisively deemed 'the main stream media', he received far more media attention than any other candidate. According to MediaQuant, a firm that tracks media coverage of candidates, Trump received nearly five billion dollars of free media attention compared to Clinton's three million (Harris, 2016). Scholars have also noted that the Trump campaign was innovative with its use of social media (Albright, 2016; Beckett, 2016). An article in Wired Magazine went as far as to say that sites like Facebook and Twitter won Trump the presidency (Lapowsky, 2016). The same article noted that "social media was Trump's primary communication channel." In a conversation with CBS's 60 Minutes (2016) , Trump himself said that Twitter and Facebook were key to his victory. Bots, according to informants, pushed these successes.

The numbers from the Wesleyan Media Project's (Franklin Fowler et al., 2016) report on campaign spending suggest that, as with the polls, the metrics by which advertising agencies seek insight into political wins proved to be misleading when it came to an actual outcome. Television advertising seemed to have very little bearing on success, Clinton spent $258 million to Trump's $100 million. On local cable Trump had less than a one percent market share. Clinton even dominated digital ads (desktop, display, pre-roll) had 73% share of nationally focused digital ads (desktop, display, pre-roll), Trump 27 Percent.

Social media's affordances for democracy, for enhancing the ability for communication and organization, have long been discussed by scholars concerned with politics and media whether for networked communication amongst online media producers, communication amongst democratic activists, or to disrupt political conversations amongst elites (Benkler, 2006; Howard & Hussain, 2013; Owen, 2015). More recently, there has been a normalizing pattern on sites like Facebook and Twitter. That is, political elites have figured out how to harness social media to exert power and control (Karpf, 2012). Donald Trump used one digital tool above all others to circumvent the need for traditional political advertising.

As one informant, a conservative social media expert named Clint, put it: "Trump used Twitter as a megaphone, as a tool to get his campaign message heard above all others." However, suggesting that the Trump campaign's success in harnessing social media, an emergent version of political normalization online, won him the presidency is off the

mark. In fact, a somewhat oppositional phenomenon, the democratization of propaganda, was also key to his success. Together, the campaigns creative use of social media and supporters' use of large scale social automation allowed the agenda of the media to be set in favor of Trump.

F. Democratizing Propaganda over Social Media

Discussions about the Trump campaign's attempts to speak over all other news, what Clint called "megaphoning," became a clear theme in interviews (interview, March 31 2016; interview, November 12, 2016). For instance, another research informant, Al, echoed Clint's claims of this amplified communication tactic (interview, November 12, 2016). Al was and is a high-ranking member of the Republican Party apparatus. Al told me that the campaigns he had worked on treated "digital" (online marketing) like the "wild west." He said, "Anything goes as long as your candidate is getting the most attention."

Generally speaking, social media bots play a fairly heavy-handed role in amplifying political messages. The idea behind political botnets is one of numbers, if one account makes a splash with a message than 1,000 bot-driven accounts make a flood. Armies of bots pretending to be human, what some call sock-puppet accounts, computationally and automatically extend the ability of the deploying party to spread messages on sites like Twitter. Political botnets are no exception. During the 2016 election I catalogued numerous occurrences of bots being used to drive up traffic around a particular event or idea.

For instance, at the height of Pizzagate, the conspiracy that linked the Clinton campaign to an alleged human trafficking and child abuse ring, automated shell accounts rampantly spread memes putting Clinton campaign Chair John Podesta and the candidate herself at the center of the fabricated controversy. A disproportionate number of the accounts generating traffic on Pizzagate appeared to originate in Cyprus, the Czech Republic, and Vietnam. According to the Washington Post, "as the bots joined ordinary Twitter users in pushing out Pizzagate-related rumors, the notion spread like wildfire" (Fisher et al., 2017). Pro-Clinton bots also spread attacks on Donald Trump, though they were about 1/5 as active as the pro-Trump bots during key election events (Howard et al., 2016).

One example of bots amplifying political messages during the campaign stands out. In April 2016, conservative political strategist Patrick Ruffini, webmaster of the 2004 Bush/Cheney campaign and former eCampaign Director of the RNC, sent out a series of tweets suggesting that bots were being used to attack Ted Cruz. The Daily Caller and the National Review quickly picked up the story, both suggesting that the bots were potentially part of a broader network of 'fake' Trump Twitter traffic.

Ruffini made a spreadsheet of nearly 500 allegedly automated accounts, many of which were deleted or became inactive just after he publically shared the list. Most of the accounts had no followers, copied one another's messages, and sent out advertisements alongside political content. Ruffini found that they were also being used to support the Trump campaign. The strategist noted that the bots sent out 400,000 messages about

Trump, and nearly 2 million tweets total, over the course of a month. The same accounts retweeted Dan Scavino, Trump's social media director, nearly 15,000 times.

Figure 3: Patrick Ruffini Outlines Bot Attacks Against Ted Cruz

Source: Author's screenshot 4/8/16.

Note: Ruffini send out several messages on Twitter about the accounts in question. He also compiled a spreadsheet of accounts, complete with metrics and links, and shared this on the platform.

Ruffini's main issue with the accounts was that they were urging those who had received Ted Cruz campaign robocalls to report him to the Federal Communications Commission for violating political media regulations. So, a group of automated Twitter accounts were being deployed to mobilize voters against automated campaign phone calls. This tactic, of using bots to make assertions about campaign law, is something I had not seen in two previous years of research about political bot use in other countries. Also interesting was

the fact that this was a group of pro-Trump Republican bots being used to attack Ted Cruz, another Republican. In an interview with Politico Ruffini said, "A lot of these unsavory tactics that you would see in international elections are being imported to the U.S." (Schreckinger, 2016b). He also noted, "there is very clearly now a very conscious strategy to try to delegitimize opposition to Trump."

Figure 4: Twitter Bot Pushes FCC Complaint Against Ted Cruz

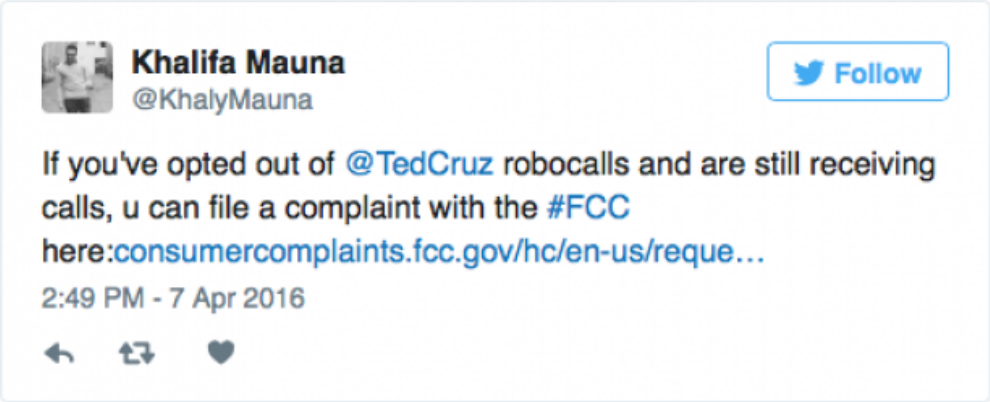

Source: Author's screencap 4/8/16.

Note: A screencap of the tweet copied and sent out by hundreds of purpose-built bot accounts. Twitter has since suspended this particular account, along with most others associated with the attack.

Ruffini's allegations, and his efforts to catalogue information about the accounts in question, provided reason for further examination into party, or candidate, focused bots. There is evidence that U.S. political actors have previously used bots in attempts to manipulate public opinion (Ratkiewicz et al., 2011; Metaxas & Mustafaraj, 2012). During the 2012 election cycle Mitt Romney's campaign was accused of buying thousands of followers on Twitter in a bid to seem more popular (Coldewey, 2012). In 2010 researchers discovered a botnet purpose built to attack Martha Coakley, the former Massachusetts Attorney General, by alleging she was anti-Catholic (Mustafaraj &

Metaxas, 2010). At the time Coakley was in a tight race with Scott Brown in the special election to fill Ted Kennedy's senate seat. Brown eventually won the race. In these cases, bots were used to support U.S. political candidates, and even to attack opposition. Was this a common campaign tactic, however?

The same week that the anti-Cruz botnet was launched I managed to make contact with a well-placed member of the Republican Party. The informant, Jane, had worked on several high profile political campaigns and was, at the time, employed by the Republican National Committee. When I asked her if she had seen campaigns use social media bots before she answered bluntly, "Yes, absolutely. It is a common tactic, in both presidential campaigns and lower down the ladder." She was, however, skeptical that using bots to boost candidates was actually effective. In fact, Jane told me that doing so was, in her opinion, a waste of money and more a distraction than a benefit. She said, "likes and retweets don't equal votes." That said, she told me that in her experience digital teams treated online strategy in a fairly ad hoc way. "We will throw anything against the wall and see what sticks," Jane said, "bots are one tactic among many, and they aren't illegal."

There are two clear take away points from my interview with Jane. The first is that, despite her own ambivalence about the efficacy of political bots, she openly admitted that bots were regularly used in Republican campaigns at all levels of governance. Second, she was emphatic that digital campaign teams, again at all levels, commonly made use of a variety of tactics and treated the online space as a frontier for testing new marketing methods. This picture, one of experimentation and spontaneity, stands in stark contrast to

the one painted later in the in the campaign by groups like Cambridge Analytica. Jane was not alone in this assessment, several other informants who worked in digital campaign contracting echoed her claims. Some went further, saying that claims of psychographic or psychometric targeting were largely exaggerated and that it was a clear campaign message and boots on the ground led to votes, not fancy computational tactics.

The democratization of online propaganda is explored more deeply in chapter Five, which looks at bot usage by digital entrepreneurs. This phenomenon has occurred partially because the use of political bots, and infrastructural elements of sites like Twitter such as an open application programming interface (API), made it possible for average citizens, people outside of Washington or the professional campaign apparatuses, to amplify their own viewpoints online. The reach and sheer numerical strength of bot driven communication allowed anyone with some coding knowledge, or connections to groups using automated software, to create their own propaganda network.

The question of whether campaigns themselves used political bots to spread "fake news" or in collusion with foreign powers was, and continues to be, a smoking gun issue in US politics. However, the democratization of online propaganda—while less provocative or comprehensible due to its amorphousness and separation from the centralized power in Washington D.C.—was none the less an incredibly important part of how computational propaganda and political bots were used in the 2016 U.S. election. While government departments, academics, and journalists continue to search for evidence that campaigns

used various underhanded digital means to manipulate public opinion, they tend to ignore the fact that anyone can launch a bot—or spread fake news—online. It was these citizen-built bots that accounted for the largest spread of propaganda, false information, and political attacks during the 2016 election. They were built to be explicitly partisan, to trumpet news and opinions in favor of the candidates or to promulgate attacks against them.

Chapter 4

Information Radiators:

Journalism, Bots, and the 2016 Campaign

A. Introduction

In 2016, for the first time since 1944, the presidential candidates of the two primary U.S. political parties called the state and city of New York home. The Clinton campaign based its operations out of an office building in downtown Brooklyn. Trump's core team worked from his eponymous tower in Mid-Town Manhattan. New York City is also a hub for global communication. It is the base for many of the world's leading news outlets, magazines, and publishers. The offices of three of top five most circulated U.S. daily newspapers are located in the city, as are the headquarters of NBC, CBS, and ABC and a multitude of other influential media organizations. More so than any other U.S. city, and perhaps any city in the world, New York is the center of the contemporary information environment and some of the biggest global media regimes (Williams & Carpini, 2011). Unsurprisingly, a huge proportion of reporting on the 2016 U.S. presidential race came out of the city. Many innovative journalism teams, who built and leveraged online technology in new ways during the contest, were stationed there.

Though Trump and Clinton dominated many of the headlines in the year leading up to election day, they regularly shared news space with an unlikely online tool—the bot. For the technology world in general, and Silicon Valley specifically, 2016 was the year of the

bot. Executives from the world's biggest tech firms, including Microsoft CEO Satya Nadella, situated chat bots as the logical technological successor to the app (Bruner, 2016). At the company's World Partner Conference 2016 Nadella said that "the evolution of user behavior [in] computing will naturally involve the use of chatbots as the underlying basis for communication between user and information" (Anderson, 2016).

In the midst of the excitement around bots, an array of startup companies built their business on using social bots in a variety of novel ways. The chat and collaboration platform Slack, launched just three years before the 2016 election, allocated $80 million to a fund geared mostly towards supporting bot driven start-ups. This included millions of dollars invested in 11 bot making companies by the platform in the first six months of 2016 (Cifuentes, 2016). Journalists, who had been tinkering with bots on Twitter and Facebook since the early days of those platforms, built several novel bots to parse and share information in 2016—many specifically related to the election. Diana, a reporter and designer at a newspaper in New York, told me that bots were crucial tools for journalists hoping to understand the profusion of data online. She said that bots could—and were—being built by journalists at her organization and others like it in order to automatically "communicate with people around the U.S. about the election" (interview, March 30, 2016).

Amidst the excitement about bots of all varieties, and with the pivotal 2016 U.S. election as a test case, numerous U.S. news organizations quite suddenly began scaling efforts to launch journalism bots: software tools used by journalists to automatically parse

information in large databases, identify online trends, and communicate with readers. Journalism bots were deployed during the 2016 election on the major traditional social media sites: Facebook, Twitter, and Youtube. They were also used by newsmakers on emergent social media and closed network chat apps like Slack, Facebook Messenger, and WhatsApp. Legacy newspapers such as the *New York Times* and the *Washington Post* as well as newer web-based publications such as *Buzzfeed* and *Quartz* experimented with using bots to both report and gather information upon the presidential race (Hazard Owen, 2016; WashPostPR, 2016; Mullin, 2016; Wang, 2016). These, and other, news outlets used Twitter bots to share breaking news, Slack bots to relay questions from readers to journalists, and Facebook Messenger bots to provide the latest poll numbers.

At the same time, political actors hoping to gain the upper hand in the contentious race leveraged social media bots to control political communication rather than facilitate it. Manipulative political bots, used to harass journalists and confuse the flow of information, were used alongside journalism bots during the 2016 U.S. election. These tools maintained functions markedly similar to their reporting-oriented cousins. They were built to automatically communicate, but not for the sake of clarifying information or facilitating knowledge sharing. Rather, they made use of anonymity and large numbers of sock-puppet social media accounts to give astroturf support to stymie the democratic communication process.

Political bots were among the primary communication technologies used over social media during the 2016 contest to spread mis- and dis- information (Shao et al., 2017).

They were used to amplify troll attacks against reporters and political opposition (Guilbeault & Woolley, 2016). Analysis from Monaco and Nyss (2018) on the rise of state sponsored trolling notes that the now global trend of "fake news is often only one element of a broader politically-motivated attack on the credibility and courage of dissenting voices: journalists, opposition politicians and activists." Bot driven tactics for manipulating public opinion and controlling communication reached new heights during the election in 2016 and many of the tactics used during this contest have been used in more recent political events around the globe.

Bots were critical to the spread of what has now become popularly known at "fake news", a term itself the cause of much partisan disagreement and public confusion. On the day of the election a software engineer named George, who had experience building applications on social media, told me that social media bots could just as easily be built to spread false information as they could to spread facts (interview, Nov. 8, 2016). In the days preceding the election, bots were used on Twitter to amplify hashtags related to the Trump and Clinton campaigns (Kollanyi et al., 2016). In the week preceding the election, over five times as many bots tweeted in support of hashtags associated with the Trump campaign (#MAGA, #Trumptrain etc.) than those associated with Clinton (#Imwithher, #Clinton2016, etc.). Political bot makers launched massive bot armies, known as botnets, to alter communication over social media about the presidential race. At certain pivotal moments during the election, bots generated as much as one fifth of the conversation about the race on Twitter (Bessi & Ferrara, 2016a). According to those I spoke to during

my field work, bot builders used similar botnets to circulate conspiracies about the candidates and to exacerbate the partisan divide (interview, Sept. 29, 2017).

B. Overview

From February 2016 to November 2016 I made a series of trips around the United States to spend time with people using bots to facilitate—and interrupt—the communication of the news. I spoke to data journalists specifically reporting on stories related to the presidential campaign. I interviewed and hung out with investigative reporters and digital support staff who worked for major publications based in New York, Los Angeles, Chicago, Seattle, and Detroit. I spent time with organizations building new digital tools for journalists. I also worked with groups attempting to grapple with the effects of disinformation upon the larger media environment. In short, had numerous conversations with people building, launching, and interacting with bots in and around journalism during the 2016 race.

After the election, and throughout most of 2017, I continued to travel for research. Informants made it clear that 2016's electoral result—and the marked rise in concern about "fake news" and manipulative political bot usage—was the cause for deep introspection amongst the news making community (interview, Feb. 2, 2017; interview, Feb. 3, 2017) It was also a catalyst for innovation, a wake-up call for fact checking organizations, data journalists, and tech companies working in and around trust and verification. These groups are still grappling with ways to use bots to support news production just as they are attempting to avoid being the target of bots working to silence

them. It is also clear that the usage of bots both by and against journalists is spreading to new digital domains—from major web 2.0 platforms such as Facebook and YouTube to newer applications such as WhatsApp and Telegram (Chesney et al., 2018).

This chapter uses data from interviews and participant observation to illuminate the role of bots in and around the 2016 U.S. presidential election. I compare what I learned from journalists in New York City and Washington D.C., near the heart of the race, to the perspectives I gathered from their colleagues during field work in a variety of other cities and towns around the country. I examine the ways in which news outlets used social automation in attempts to both gather data for stories and streamline communication between newsrooms and readers. I also unpack the role of bots in a diverse series of efforts to stymie professional journalism and spread junk news. In order to understand the role of bots in news gathering and production during the contest I examine how journalists used bots in efforts to innovate information sharing, but also how variety of political and commercial actors used bots to muddy understandings of information.

How, in particular, did professional journalists and amateur news makers use bots—and how were bots used against these communities—during the 2016 contest? How were bots used to facilitate the spread of professionally produced news? How did bots aid in the spread of misinformation, inadvertently spread false information, and disinformation, purposefully spread false information, during this pivotal political event? Why and how did bot builders prioritize particular uses for bots spreading information in and around

journalism? All of these questions are central to this chapter and my research into the use of bots by and against journalists during the 2016 contest.

C. Understanding Journalism Bots

Though the use of bots to facilitate journalism via social media is a relatively new phenomenon, pertaining mostly to the rise of Twitter, YouTube, and Facebook in the early 2000s, their use during the 2016 U.S. case was significant in several ways. First, journalists used bots to communicate in new ways during the presidential contest. At the *New York Times*' R & D Lab the team used bots to connect readers to reporters so that they could ask questions—and provide story leads—about the election. *Quartz*, a newer online news magazine owned by the same parent company as *The Atlantic,* launched its own "bot studio" in November 2016, just weeks after the election. Zach Seward, Quartz's executive editor and VP of product prioritized Slack—not Facebook or Twitter—as a primary area for innovation in journalism bot building (S. Wang, 2016). He said, "if we're going to make a tool or a set of automated tools for journalists I imagine we'd build something for Slack, with so many newsrooms operating largely on Slack now" (Wang, n.p.).

The usage of bots over Slack, for companies like Quartz and others around the globe, have two primary communicative functions. First, they are useful in connecting journalists to content relevant to particular stories. A bot might, for instance, be built to

fetch other writing related to a topic. Diana, a journalist and coder, told me that "in an age where data leaks and whistle blowers are as technologically oriented as they are revealing, bots can crawl through thousands of pages online—searching for names, phrases, words, or numbers" (interview, March, 30 2016). Bots are also useful in connecting newsrooms and publics, acting as social scaffolding for groups that might otherwise have difficulty connecting. As Diana put it, readers could interact with the bots that she and her team had built and launched over Slack by "ask[ing] a question, or [telling] a fact, that would then be relayed the media organizations reporting team to then be either answered or analyzed."

Researchers have examined journalism bots capacities to search for information and communicate with readers. Some have used these automated digital tools to cursorily problematize the idea that communications necessarily exists between two or more people, arguing instead that tools like bots play a key role in news sharing online (Larsson & Hallvard, 2015). Others have argued that they can produce valuable news during political events and as such are likely to play a key role in the future of political reporting (Gonzales & González, 2017). Lokot and Diakopoulos (2016) propose a typology of "news bots" in order to guide intent, utility, and functionality of bots constructed by future designers and reporting teams. While each of these authors note the limits of robot journalists—especially in the areas of automated commentary, opinion writing, algorithmic transparency, and general accountability—they tend towards limiting their arguments about the positive functions of news or journalism bots with suggestions that these automated tools could prove useful for niche journalism or local reporting.

The analysis of Lokot and Diakopoulos (2016) primarily focuses on design elements of a sample of extant news bots on Twitter. They examine the various journalistic functions of these accounts but do not interrogate the intent of those who build the accounts. While they make it clear that news bots could change the modern media environment, their exploration of several accounts is concerned with their function in generating articles and reporting rather than with either the power dynamics inherent in the deployment and use of such tools or the ways that bots can be constructed to critique the power of others. They discuss problems associated with opacity of the algorithms that drive news bots but leave room for a larger discussion about the people who construct those algorithms: what cultural values they encode into that software, and the function and of the resultant bots during political crises and elections.

Pan, an educator and software developer, was particularly reflexive about the journalism and commentary he and his colleagues constructed (interview, Jan. 13, 2017). He, in contrast to many journalists I spoke to, made it clear that his bots were specifically build to make political commentary rather than to produce objective updates about current events. One of his creations, for instance, accessed facts from both past and present to critique Trump's policies. Pan told me he though a Twitter bot might be a good "fairly easy, automated, and functional" way of posting criticism. Rather than a news reporter, he considered himself "a humanities research hound," who combined "intention, politics, and knowledge with research and the desire to do good".

Pan's creations, inarguably a form of opinion oriented journalism bots, were specifically designed to address issues he saw in U.S. politics. In fact, he made it bluntly clear that his computational work enabled him to critique over social media in a way—and to a crowd—he would not be able to access through offline means. He said he always tells his students, "computers are stupid, they do exactly what humans tell them to do." Pan's comments suggest that only definitional issues around Lokot and Diapoulous' (2016) conception of news bots are cause for their suggestion that such tools are limited in their ability to produce commentary or opinion work. Rather, discussions with bot builders and designers reveal that while the journalism bots they build might not produce traditional opinion articles in the vein of the *New York Times* or the *Washington Post*, they certainly use current events, news, and opinion to generate a form of critical reporting.

D. Bots and the 2016 U.S. Election: Bots for and against Journalism

The usage of bots by and against journalists during the 2016 US presidential election was varied in terms of both the intentions behind it and success in generating meaningful change. Journalism bots and political bots exist on a spectrum of uses ranging from democratically and communicatively constructive (journalism bots) to civically and communicatively repressive (political bots). It is how these automated tools are used that leads to features that add to or detract from the open flow of information. One informant, a data journalist named Reke, was blunt in her assessment of the field of bot usage around journalism. She was emphatic that more "tools are needed to work towards identifying and disarming what she called "bad" bots and supporting "good" ones (interview, Sept. 27, 2016).

The use of journalism bots for the purposes of information parsing and news dissemination reached unprecedented levels during the 2016 U.S. election. As part of the Computational Propaganda Research team at the University of Oxford I engaged in analysis of general social media bot usage during numerous prior political events: the governmental crackdown on free speech and journalism in 2014 and 2015 Turkey, similar events in Ecuador in the summer of 2015, the Brexit referendum in the UK in 2016, and several others. In none of these events were journalists able to effectively harness the journalistic power of bots in the way their U.S. colleagues did in 2016.

Multiple U.S. publications built and deployed social bots specifically tasked with sharing news and engaging with readers on subjects related to the election. Shelon, an editor at an online news site, told me that social media bots held tremendous potential for allowing journalists more time for important tasks like writing, interviews, and investigation of stories (interview, Feb. 2, 2016). She told me that she had never seen journalists experiment with social media bots in the ways they were in 2016. For her, the most interesting function of the journalism bot was as a proxy of the journalist. She told me it would could be tasked with doing the more mundane work of automated online "heat seeking". Shelon told me that these activities included cataloguing the latest trends on different platforms, searching spreadsheets for important information, or updating social media feeds with new stories.

At the same time bots were being used for creative new reporting purposes, they were also being used to attack members of the journalism community, particular news outlets, and the general tenants of the fourth estate. Shelon (interview, Feb. 2, 2016) told me that she had colleagues both in the U.S. and abroad that had been on the receiving end of bot-driven troll attacks. These automated offensives against certain segments of U.S. journalism have continued post-election. Political bots were used alongside digital sock puppet accounts—human operated accounts using fake or masked identities to mimic real users—as well as highly coordinated human-run accounts in efforts to levy computational propaganda against journalists. Amplified human-automaton armies of astroturf constituents were successful in promoting and spreading narratives counter to professionally produced news. Much of this content contained what has come to be known as "fake news."

Bots were critical in pushing disinformation, or fake news, during the 2016 election. Many media outlets have argued that this emergent form of computational propaganda helped to secure Donald Trump's eventual victory (Read, 2016; Parkinson, 2016). The term fake news has itself become contested, with some arguing that it "has been weaponized by politicians over [2017] to discredit media reports that they dislike" (Funke, 2017). With this in mind, I use Howard et al.'s (2017) concept of "junk" news in lieu of fake news. They define junk news as news-like content that is "extremist, sensationalist, conspiratorial, masked commentary, fake" or otherwise suspect. They argue that "social media platforms have "served significant volumes […] of junk news at sensitive moments in public life, though most platforms have revealed little about how

much of this content there is or what its impact on users might be" (p. 1). Furthermore, they argue that bots are a "pernicious means of spreading junk news over social media to friends and family" (p. 1).

Political bots were instrumental in helping several major junk news stories to go viral in 2016 (Fisher, Cox, & Hermann, 2016; Robb, 2017). Perhaps the most popularly known manufactured scandals propelled—in part—by bots was "Pizzagate". This conspiracy theory, spread through social media from Reddit to Twitter, suggested that Hillary Clinton and other major Democratic figures were engaged in an illicit sex ring housed in a pizza parlor in Washington D.C. Pizzagate was propelled by automation and human sharing over multiple social media platforms, picked up and covered widely in the U.S. and international media, and eventually led to serious offline consequences. Tomas, a freelance journalist, told me that he when he began to dig into the pizzagate story he found that many of the accounts central to spreading the story on Twitter messaged hundreds of times a day about the same content, lacked normal profile information, and were followed by large numbers of similar accounts. All of these are characteristics of highly automated, or bot, accounts (interview, Sept. 29, 2017). Jonathan Albright, a professor at Columbia, described the role of social media automation in Pizzagate: "What bots are doing is really getting this thing trending on Twitter, [they] are providing the online crowds that are providing legitimacy" (Fisher et al.).

Field research revealed several other themes in the usage of bots for spreading political propaganda during the contest. Informants provided detailed accounts of coordinated,

thousands of accounts strong, bot attacks upon specific reporters or reporting communities (interview, Feb. 2, 2016; interview, Dec. 1, 2016). In one case, a botnet was used to target multiple Jewish journalists in an organized campaign using the hashtag #thelist (Fleishman, 2016). According to one researcher I spoke to, who has spoken to several of the journalists attacked, #thelist was effective because most of those targeted chose not to report that they had been attacked (interview, Oct. 27 2016). Many also stopped reporting about the alt-right—who they believed was behind the attack—for a period, she said.

With these examples in mind, I argue that these types of political bot driven attack is an emergent means of digital political communication for control. Bots are specifically effective in amplifying junk news articles that might not otherwise gain purchase during an election. They are used to magnify attempts to silence and harass journalists reporting on contentious political events. Powerful political actors now use social media bots in attempts to suppress particular social media narratives by causing chilling effect amongst journalists and political opposition. In the legal context chilling effect is the result of an action that "inhibits the exercise of first amendment rights" (Schauer, 1978). To achieve these ends, groups relying on the anonymity provided by the digital sphere deploy political bots to overwhelm journalists' social media accounts with threats, vitriol, or spam. The goal is to force them offline—either from fear or frustration—to censor them. Most cases of this happen when the journalist under attack is reporting on politically sensitive stories.

Conversely, bots can also be used to facilitate and even improve the journalist's ability to report. Early in my time studying the makers and trackers of bots I spoke to AJ, a reporter who had built several Twitter bots to share news. He called these journalism bots "information radiators" (interview, March, 31, 2015). My discussion with AJ, and the evidence I saw in the field, suggests that an information radiator is best defined as a social bot used on behalf of a journalist to more effectively—even continuously—report on stories using social media.

What, AJ mused, should a journalist do with all the other information that came across their desk? How could the journalist get the masses of information relevant to the public, congressional voting records for instance, out to citizens without writing stories? This is where, for AJ, bots could be used to facilitate updates or report on simple stories. Information radiators, as well as other types of journalism bots, can be used as social scaffolding, or a reporting prosthesis, for a variety of news making communities. Concurrently, they can also be used to bolster online echo chambers, sow bipartisan division, and to confuse reports on current events.

The bot as an information radiator highlights the difference in ways that journalists used bots and how many candidate/campaign or digital constituent bot builders used them. At its heart this is difference is about how bots can be used to support democracy versus how they are used to promote consensus—or the illusion of consensus. When informants who worked as journalists said they built bots to extend their ability to report on facts and make information more transparent they often spoke about this usage in terms of how this

supported democracy (interview, March, 31, 2015). AJ consistently returned to the idea that his Twitter bots were used to clarify elements of American politics that would otherwise remain obscure—how a politician voted on a bill or what they said about a particularly controversial topic. He said that his bots held democratically elected officials accountable for their actions. This use stands in contrast to the ways that several digital political consultants spoke about using bots. They discussed how armies of bots were useful in making an idea seem popular. They wanted to manufacture consensus and control voters rather than radiate information for democracy.

E. Journalism Bots: Utility, Novelty, and Complexity

On one of my trips New York City during the 2016 U.S. presidential election, I was able to visit and interview a journalist and designer named Leah at the major U.S. news publication where she worked. I was interested in Leah's perspective on the use of bots in and around journalism because I knew that she and her team had been experimenting with building interactive journalism bots for use during the 2016 election. The bots in question were made to enable better one-to-one interaction between readers and reporters. Leah told me that she had been working with bots for many years and that she was interested in them for their "botness" as much as she was interested in them for their ability to "mimic people" (interview, March 29, 2017). She said that she believed that the most special aspect of the social bot is wonder: "bots are at their best and worst," she said, "when they do things that we don't expect."

Leah wanted to use bots to push the boundaries of journalistic practice (interview, March 29, 2017). She made it clear that she saw myriad uses for bots as tools for journalists searching or reporting on information. More interesting to her, however, was that their social functions and automation—alongside their relative freedom in interacting with ecosystems online—oftentimes pushed the boundaries of journalistic expectation. At times, bots that she had already built and launched online revealed novel uses that she and her reporting team had not considered when scoping their construction. A bot could, for instance, be used to simultaneously gather trending stories on the 2016 election, amalgamate information from the latest polls, and share that information with both reporters and readers through social media. At times, however, a bot set to do these tasks would surface bits of information that the reporters did not expect: discrepancies in polling data, newsworthy stories shared with the bot via a reader on Slack, or a penchant for unintentional news humor.

Leah and her team were experimenting with the use of journalism bots on closed network chat apps such as Facebook Messenger, WhatsApp, and Slack (interview, March 29, 2017). One bot, which they released publically for a trial period, allowed for better communication between journalists and readers. This bot was used as an intermediary that allowed readers to ask questions, or even provide leads, to journalists. The bot was also, however, made to fetch basic information for readers: the latest headlines on the race, recent polls, or information on candidates' official political stances. The bot was ultimately, however, not successful. Leah told me that many people found the bot's multiple functions confusing. What is more, many thought that the way that it

communicated was annoying and, at times, spammy. In retrospect, she said, the bot should have been made to be more nuanced, to be able to be programmed to the user's desire. If the user only wanted to see poll updates, then the bot should have been able to do just that. As it was, it tried to do too much and this resulted in an information overload, and confusion, for users. The functions of a bot, she made clear, do not need to be complex to be successful. In fact, she believed this bot would have been much more successful if it had discrete, personalized, uses.

It is common for people to regard social media bots, and their conversation enabled cousins chat bots, with a degree of curiosity because they are both automated and socially capable. Diana told me that she thought there was something uncanny about communicating with a disembodied personality (interview, March 30, 2016). Harry, an amateur bot builder, specifically built his first bot to confront humans with the alien nature of algorithms and computers (interview, Feb. 2, 2016). He told me that his work seeks to challenge passive acceptance of automated technologies and was adamant that people should think about the weird nature of bots as extensions of computation. When bots are programmed to interact with other users online their automation means that they can retrieve a relevant response from a data base of phrases and write back almost as soon as their conversation partner enters a phrase. This instantaneousness can be unsettling. Moreover, advances in machine learning allow chat bots to more readily learn from their environment and to use what they pick up in their conversations. Tay, Microsoft's failed Twitter chat bot experiment, was built to learn from other users

(Vincent, 2016). Without safe guards, this quickly let to the bot regurgitating hateful phrases after platform trolls realized that she had a simple "repeat after me" mechanism.

The majority of bots on Twitter, however, are not very complicated. Leah noted that, in order to achieve the purposes their builders set out for them, a bot does not necessarily need to be complex (interview, March 29, 2017). This is also true for political bots used to artificially amplify an idea, person, or junk news article. For instance, if a client hires a botnet—or large group of bots—to make their social media account look popular than the easiest way to achieve this is to have the bots follow the account and positively interact with the account in a simple, and relatively passive, way. The bots can simply like or re-message the account's content.

Social media bots get used in and around the field of journalism for a variety of purposes. Bot building journalists revealed a plethora of ways they could harness bots for civic good. They described making and deploying bots that could automatically crawl massive data sets and others that could use social media to publically communicate important findings. Bots that communicate information over a platform like Twitter fall into the category of AJ's information radiators (interview, March 31, 2015). They make it possible for reporters to share salient details from massive swaths of information from sources like Wikileaks, which would otherwise be nearly impossible to effectively comb manually. Other bots can be launched over social media to monitor and report on potentially important online trends or findings. These are what Shelon described as "heat seeking" bots (interview, Feb. 2, 2016).

In some ways, bots act as a sort of prosthesis or scaffolding for the data minded journalist. They are tools that extend and support reporters' ability to do their job. Some journalists and researchers have, however, also made the argument that bots could work to replace human journalists (Lokot & Diakopoulos, 2016). Indeed, tools like the LA Times Quake Bot can automatically generate and post stories (Walker, 2014). In 2016, Harvard University's Nieman Journalism Lab argued that there will be a large-scale shift towards the "botification" of the news in coming years (Barot, 2016). Other news outlets, with more unease about automated reporting, have suggested that journalism may be the latest industry to be under threat by automation—because of article writing bots—or that algorithms may "kill" journalism (Goichman, 2017; Keohane, 2017).

Some of the bot builders and journalists I spoke to were less than worried about bots supplanting journalists. In fact, AJ argued that bots could do mundane tasks while journalists could focus on in-depth reporting or other aspects of news making that most excited them (interview, March 31, 2015). Leah told me that most journalists do not want to report on little league games or minor earthquakes (interview, March 29, 2017). She said that she was fine with bots being used to write simple articles wherein the bot could use a basic formula. AJ detailed building many such bots. One tracked congressional stances on gun control. This Twitter bot, AJ said, was simply programmed to constantly tweet politicians' tweets on the issue—making them available for the public. Anytime a politician tweeted something related to guns or gun control the bot would gather the information and retweet it. It was inspired by impending measures related to gun control:

the assault weapons ban and a widely-discussed bill on background checks. AJ released the source code and described the coding process for this bot publically.

AJ also described a news bots inspired by his gun control bot (interview, March 31, 2015). This particular bot was broader in scope, he told me, because it was both an information radiator and, though he didn't use this specific term, a heat seeker. It was built to follow people and gather information on their tweets. The gathering apparatus could be tuned to any subject and could then be built to tweet information on the subject. He told me that he could imagine multiple other journalistic uses for this type of bot. It could gather information from all over Twitter—sentiments from the public to statements from public figures—and reveal particular ideas, feelings or beliefs that might otherwise go unnoticed.

The various uses of journalism bots underscore the fact that bots are communication tools. They can be used to enhance the capabilities of journalists, just as they can be used to enhance those of political campaigns and their supporters or of digital constituents. The most salient characteristic of bots, as they are used in the field of journalism or elsewhere, is their automation. Bots allow journalists, and often those who distinguish themselves as data journalists, the ability to automatically sift through terabytes of information that would otherwise take human teams years to excavate. On the forward-facing side of social media, social bots can act as a proxy for their creators.

F. Why Journalism Bots?

Journalism bots can gather, parse and communicate information online. Sometimes they are built in house at a newspaper by a digital support team. *Quartz*'s bot studio, with a mandate to "experiment with applications of bots, AI, and related technologies for journalism on new platforms" (Seward, 2016), build in-house bots across multiple social media platforms. Other journalism bots are constructed by technology firms, think tanks, and news savvy individuals. Many journalism bots are attached to front-end facing profiles on social media. Hoaxbot, for instance, is a "friendly Twitter bot that replies to people who tweet links to fake news articles" (König, 2016). The *New York Times* continues to experiment with bots that generate a more one-to-one experience between reporters and readers (Phelps, 2017). The *NYT* Political Bot, for instance was constructed during the 2016 U.S. election to share key race updates over Facebook (Lichterman, 2016).

AJ told me that bots can be used by journalists and other democratically minded individuals or groups "to keep politicians and public figures honest about political stances and actions but they also tweet news related humor" (interview, March 31, 2015). Harry elaborated on this same point (interview, Feb. 6, 2016). He told me that bots could automatically critique politics and politicians while their maker slept.

There are also journalism bots that prioritize automated fetching, cataloguing, and storing mechanisms. Quackbot, a Slack bot project by Quartz and the document analysis company DocumentCloud, "performs tasks useful to reporters, editors, and news

producers" that range from "grabbing screenshots of webpages to pointing out clichés" (Keefe, 2017). Buzzbot, a project of Buzzfeed's Open Lab, combines both a news gathering mechanism with a social media reporting function. Launched during the 2016 Republican National Convention, this journalism bot was conceived as a tool for helping small groups of journalists cover massive events more successfully by inviting other people to participate in coverage. The team that constructed it sees it as "the vanguard of bot reporting" a new form of journalism "which BuzzFeed Editor-in-Chief Ben Smith considers an interesting category that might free up journalists to do other stuff" (Mullin, 2016).

Bots used to attack journalists or spread fake news are purposefully not defined here as journalism bots. The former are a version of political bots deployed to monitor "hotly contested topic areas" in order to "harass and attack individuals and organizations in an attempt to push them out of the conversation" (DiResta et al., 2017). The latter are also political bots, most generally amplifiers or sockpuppets. Amplifier bots are used to massively boost a particular idea—in this case a rumor or spurious news story—by sending it out at a computationally enhanced rate and/or through a botnet of thousands of shell accounts. Sockpuppets, or cyborg accounts, are run by both humans and bots and are used to seed new ideas and drive discussion. This online astroturfing mechanism can be particularly damaging to journalism and journalists because it combines human intelligence and communication skills with automation. With this, proliferators of online harassment campaigns can realistically scale attempts to promote chilling effects around

particular stories while also spreading manufactured—and often fake—counter narratives (Higgins, 2016).

Many of the journalism bots used by journalists to parse or spread news during the election communicated political data, enabled political interaction between journalists and the public, and tracked politicians and their actions. Other journalists use bots, however, to communicate news on sports, fashion, the arts, and other topics. In other words, not all journalism bots communicate about capital "P" politics. There is however, a secondary way that these automated tools are political relating to the power dynamics of how and what journalists communicate to the public.

When journalists use bots as information radiators they program these tools to share specific content. This is prioritized by the reporters or the media teams that build and launch the bots. Bots, as communication tools, are therefore are imbued with agenda setting power, to tell the public not necessarily what to think, but what to think about (McCombs & Shaw, 1972). In this sense, then, they are working on behalf of those who launch them to alter public opinion. Even when a journalist works to be objective, they cannot avoid prioritizing particular information. Nor can journalists fully avoid placing their perspective—including norms, values, or beliefs—within stories in one way or another. Journalism bots are similarly biased because they are tools constructed by journalists. In some ways, journalism bots can even obfuscate the fact that the information they communicate is value laden. Like social media algorithms, bots are

often seen as independent technical actors divorced from the intent of those who build them (Gillespie, 2014).

Because the manipulative uses of political bots have been more widely studied than more benign uses, and because they are more reported on in the media, there is a tendency to assume that all bots used around politics are nefarious. However, since the beginning of our research on the Computational Propaganda Project at the University of Washington and University of Oxford, and indeed since our team coined the term "political bot", we have noted the fact that the value given to a particular political bot is in the eye of the beholder (S. C. Woolley, 2016). We have also worked to underscore the myriad beneficial democratic uses of political bots. Worryingly, however, the democratically beneficial uses of bots by journalists in attempts to enhance the flow of information were overshadowed during the 2016 election by more nefarious uses aimed at confusing or countering said flows.

G. Bot Usage Against Journalists

Though both political bots and journalism bots were both used in attempts to communicate particular information during the 2016 U.S. election, two features set political bots apart. First, unlike the journalism bots used during contest, which news entities clearly marked as bots, the political bots used against journalists and news flows tended to impersonate real social media users or mimic the accounts of news websites. Second, the primary role of political bots used during the 2016 race was to artificially amplify junk news or cast aspersions on the professional journalism and journalists. Over

the course of my fieldwork in 2016, bot builders and journalists alike told me it was clear that large groups of political bots were being used to sow confusion and spread alternative news narratives (interview, Sept, 29, 2017; interview, Dec. 1, 2016).

Journalism bots used by professional news outlets functioned singularly, with one bot communicating information to users who had signed up for automated election updates. AJ's bots, for instance, maintained one clear bot account on Twitter that constantly reposted about topics the journalist and bot maker found newsworthy. Political bots used to attack journalists and muddy the flow of information during the election mostly existed in large networks during the contest. They were built and launched to support one another's content by liking or retweeting messages to give the illusion of popularity and manufacture trends. They magnified offensives against particular stories with the goal of drawing human support for the perspective of their creators. The botnet used to harass Jewish journalists with threats that they had been added to #thelist was once such botnet (interview, Oct. 27, 2016). At times, it can be difficult to determine where the astroturf, botnet-driven assaults,

The bot makers I spoke to about the 2016 election suggested that it was common for those initiating bot offensives to learn from peers globally (interview, May 4, 2017). The internet is a space without borders, and research shows that open source bot code comes from many countries and is constantly built upon by others and repurposed. This mirrors thoughts about the transnational spread of botmaking techniques I have heard mentioned in other conversations I have had with bot makers and bot trackers working in and

around events in Turkey, Mexico, Ecuador and Ukraine (interview, July 13, 2014; interview, Aug. 14, 2015). Bot makers working for political actors in those countries had developed similar armies of automated accounts to attack journalists and political opposition (Bradshaw & Howard, 2017; Monaco & Nyss, 2018). The bot makers who operated automated accounts during the U.S. contest launched coordinated attacks on reporters in ways similar to those in other countries during other events.

One bot builder, a software engineer named Ziggy, told me that messages from several thousand automated accounts was a much more powerful a tool for co-opting the flow of communication than one or two human run troll accounts (interview, Aug. 21, 2016). One activist and journalist told me that the bot attacks had led him to leave Twitter and to stop using social media to share his reporting (interview, Dec. 1, 2016). This brought to mind a conversation I had had with a Republican political consultant, who told me that campaigners would use any means necessary to prioritize their political platforms and perspectives (interview, Sept. 16, 2016). He, along with several other digital campaign and political marketing experts, told me that digital politics was like "the Wild West". He also told me he knew of campaigns that had made use of politics bots for a variety of goals, from boosting a candidate's follows to vilifying the opposition.

H. Conclusion and the Future of Journalism Bots

Bots were used in and around the field of journalism during the 2016 U.S. presidential election for many tasks. I came across many instances of reporters using bots for connective communication work—in attempts to benefit civil society and open

discourse—but I also came across a large number of cases where bots were used to harass journalists. Informants made it clear that bots are tools, and that it is how they are used that allows for particular effects. Journalists, for instance made it clear that bots that produced simple articles could be useful social scaffolding or prostheses for reporters who needed to focus on in depth investigative reporting. Digital campaign consultants revealed that campaigns in the U.S. had used bots in the past to amplify candidates and particular information. Junk news, which rose to prominence during the 2016 election, was spread by armies of bots built to simply retweet conspiracies and commentary veiled as news.

Journalism bots were used by teams at newspapers and digital news magazines as information radiators throughout the 2016 contest. They updated readers across multiple social media platforms with the latest poll numbers or candidates' stances on particular issues. Bots were also deployed as heat seeking devices during the race. They were used by reporters to surface newsworthy content, identify online trends, and parse large data sets. Andy, a bot builder, summarized this capacity well, "any information is available to the bot. It is able to tirelessly sort through data"(interview, March 31, 2015). Political bots were used to cause chilling effects amongst journalists. They were levied against journalists from specific social groups, or reporting on certain topics, in order to silence or frighten them from reporting. These bots were built with intent to manipulate voters into believing junk news and create band wagon effects around salacious political content.

It is clear that bots will continue to play a role in journalism during future political events. However, public sentiment at present has largely turned against bots because of the ways these tools were used to promote manipulative contents during the 2016 election and other political events around the globe. This said, automation is particularly useful is helping the news media manage large amounts of information. It is also one way for news rooms with limited resources to produce systematic articles that do not require human oversight.

At present, bots cannot replace investigative journalists. In fact, most journalism bots are very simple and are built to continuously spread the latest stories over social media. That said, even simple bots can be used to critique and comment on policy or power. With advances in machine learning, bot communication will continue to be modeled on human communication and we may see more sophisticated version of bots designed to crowd source intelligence like Microsoft's Tay. My informants, however, have pushed me to consider what bots built to look like bots—not modeled on human behavior—could do. How could they aid in the production of journalism? How could they be used to alter the flow of political communication?

Studying the ways that journalists use bots alongside the ways that these automated programs are used by campaigns and candidates and digital constituents was a challenge. Journalists often described their bot usage in ways distinctly different from informants from other groups. The builders I spoke to who were associated with candidates and campaigns or who operated solo most often spoke about the ways that bots could be used

to amplify or suppress social media content on behalf of a candidate. These were clearly, by the definition used in this project, political bots. Journalists who used bots did, at times, use these tools to automatically post new articles—but such bots were not operating in large groups of accounts mimicking real people. More often than not they were single accounts that representing a news organization that were coded to update automatically.

It's true that producers of disinformation, which many now call fake news, used armies of bots to boost their stories—and that these accounts tended to fall under the definitional umbrella of political bot because they were purposefully built to manipulate public opinion. This practice was in some sense, however, anathema to the values presented by informants who used bots in their practice of professional or amateur journalism. Bots used by journalists at regional and national newspapers, magazines, and reputable online news publications often identified themselves as bots. Their builders often stated a goal of presenting objective reporting—rather than subjective or coercive content.

Moreover, many data journalists I spoke to had a more nuanced perception of the breadth of what a bot could be used to do on social media. They used bots to automatically post politician's stances on gun control, or to give live updates on a story. Bot building informants who worked as professional reporters also spoke a great deal about the ways they could use bots to scrape social media for information. They said they used bots to surface trends, or to locate sources. They spoke about the usage of bots for

communication and information diffusion in a broader way than the other groups I worked with.

All of this has helped me to realize a key difference between how journalist bot builders and bot builders from the other two groups examined here define or think about bots. Most of my informants who were journalists, unlike campaign or digital constituent bot builders, thought of bots in terms of utility—they thought of them more purely as tools. Some hypothesized on the personified nature of their creations when I challenged them, but when it came time to use bots for reporting they were not caught up in the real or fake, human or robot, nature of bots. Campaign and candidate bot builders and digital constituent bot builders were quicker to personify their bots. They were much more likely to think of bots only as fake automated profiles—built to mimic real people. They saw the utility of their creations, but they mostly used them to stand in an online gap where they needed support of many users in order to give an idea or story the illusion of popularity.

The concept of the bot as an information radiator comes from journalist bot builders' notion of the social media bot as a tool that extends the user's ability to report on current events. Unlike the concept of the bot as a proxy of the creator, explored in depth in the chapter on digital constituents, the idea of the bot as an information radiator is more tied to the bot's utility and less to its anthropomorphism. The bot as an information radiator is a step in the direction towards a bot as a proxy for its creator but it is more focused upon the particular professional context of the bot being used to automatically report.

Future work on the ways that journalists make use of bots should consider the ways in which this practice is both implicitly and explicitly political. By focusing upon the former idea—upon how the act of using bots is in reporting is a political act because reporting itself is never truly objective and thus political—scholars will more effectively delineate the differences between journalism bots and political bots. Exploring the latter idea has a similar outcome, but also allows for clearer distinction between bot usage in the spread of information versus disinformation. This research project made it clear that this question, one focused upon the veracity of information on social media during future political events, will be closely tied to future work on the role of social media bots in political communication.

Chapter 5

Democratizing Computational Propaganda:
Bots and Digital Constituents During the 2016 Campaign

A. Introduction

One of the biggest challenges in conducting fieldwork on and around political bot building is finding and speaking to bot makers. Research shows that professional political actors—including candidates, military officials, and political consultants—around the globe have made use of political bots in attempts to manipulate public opinion, to coerce and to control (Metaxas & Mustafaraj, 2012a; Ferrara et al., 2016; Woolley & Howard, 2017). However, much less has been said about the way that average citizens deploy political bots over platforms like Facebook and Twitter for their own purposes.

Governments and politicians have made use of political bots as mechanisms for jamming communication channels of activists, for spreading propaganda on behalf of one cause or candidate or against another, or for manipulating trending algorithms (Bradshaw & Howard, 2017). How do less-resourced individuals and groups of citizens in the United States of America make use of automation over social media for the purposes of political communication? Who are these individuals and what are their goals? What does more diverse usage of political bot technology by wider portions of the population, and not just the powerful, mean for current understandings of digital disinformation?

In March 2016, I made a trip to Detroit, Michigan to do fieldwork around the open primary in that state. My main goal on this trip, beyond immersing myself in the primary scene, was to establish contact with people working directly for the various campaigns in order to understand how, or if, campaigns were making use of social bots in their digital marketing efforts. On the day before the primary I was standing in line for a Clinton rally at the Museum of African American History. During the wait, I had a conversation with a small group of Michigan State students. While discussing our reasons for attending I told them about my research on political bots and the U.S. election. It turned out that a couple of the students had more than a cursory understanding of social media bots being used in politics. In fact, they told me they had peers who had used social bots over Twitter to spread political news. The next day, while volunteering in one of the Clinton campaigns' Get Out the Vote (GOTV) call centers, another informant told me they had a friend who was using social bots on the dating app Tinder to spread information to her romantic matches about why they should vote for Bernie Sanders (interview, March 8, 2017).

The people building these bots were not working, or even volunteering, for campaigns. They were simply using their knowledge of bots and coding to facilitate their ability to engage in politics—albeit in an automated, and often amplified, capacity. This was at a time when neither bots, nor the idea of "fake" news, were at the apex of the U.S. zeitgeist. In my conversations with security experts and computer scientists in the months and years prior I had to go to lengths to explain what I meant when talking about the use of socially-oriented bots in online political communication. Most of the experts I spoke to had conceptions of bot use in these areas that were sophisticated and bound up in the

broader autonomous software use—web scrapers, spiders, crawlers—or as virally infected computers used for distributed denial of service (DDoS) attacks. These complex tools, however, are not the kinds of tools that college students use to communicate with people in their extended networks in their everyday lives. Political bots, it turns out, are.

Many interactions with "constituent" bot makers similar to those I had in Detroit made it clear that I needed a third actor category not captured within campaigns or media groups. Originally, I had planned to have a third category of political bot users encapsulated by "civil society groups". This decision was based upon my research experience around political bot building in other countries, where I had witnessed well-organized political non-profits making use of social bots. During my research around the U.S. primaries, however, it became clear that this grouping—at least as it is understood to be related to broader non-governmental organizations and institutions—did not adequately capture the lion's share of non-campaign and non-journalistic political bot making that I was seeing. In fact, focusing on civil society as one of the actor groups began, both through research and discussions with colleagues, to be too mired in challenges of definition. The description did not really fit what I was seeing and was not a discrete enough categorization. As Seligman (1995) puts it, civil society means "different things to different people" and "the resulting picture is one of great ambiguity and not a little confusion" (p. ix).

With this, and my experiences interacting with bot building individuals not connected to larger political groups or media organizations, I landed upon "digital constituents" as my

third, more specific, grouping. This final group is made of various individuals in the U.S. who are engaging in the use of political bots but who are not directly connected to campaigns or media institutions. Nor are they connected to structured non-governmental organizations (NGOs) or the like. The word "constituent" is useful in describing these individuals or small groups for two reasons related to the definition of that word: 1) they are "component[s] or parts" of the U.S. public and 2) most often "member[s] of an area which elects a representative to a legislative body" (Oxford Dictionaries, 2017). They are "digital" constituents because they are exercising their right to engage in the electoral process through free speech, here to communicate about political campaigns online.

Social media platforms such as Twitter and Facebook are transnational communication mechanisms for political communication (Segerberg & Bennett, 2011; Tufekci & Wilson, 2012). That is, their use in this regard—at least in most country cases worldwide—is not necessarily restricted by state borders. As such, people from around the globe use them to communicate about political issues with one another. Revelations from Facebook and Twitter that Russian citizens used social media and political bots to spread content during the elections have had a large impact upon the ongoing U.S. Congressional investigations about the role of disinformation during the 2016 contest (Allcott & Gentzkow, 2017). This chapter, however, does not focus on foreign individuals who used political bots to engage in topics related to the campaign. It is, rather, focused on those citizen users who worked to manipulate others from within the U.S.. The information here comes from interactions or experiences I had during fieldwork in various states with these domestic bot builders and users.

This chapter explores the ways that these digital constituents made use of social media bots for political purposes during the 2016 U.S. presidential election. I present two key ideas related to this particular variety of political bot usage: 1) the democratization of computational propaganda and 2) the bot as a proxy for the creator. I argue, for the first point, that the sociocultural differences explored in literature on the digital divide and elsewhere allow for democratization of computational propaganda in some countries and not others. Infrastructural elements of some social media platforms—including open application programing interfaces (APIs) on sites like Twitter and lowered barriers for learning to code and/or for using third party software building services such as If This Than That (IFTTT)—also allow a wide array of people in the United States to build and launch political bots.

For the second point, I suggest that social bots function as proxies, stand-in tools that are separate but connected to their makers. There is a tendency amongst scholar, journalists and the broader public to separate the social bot from its creator and, in so doing, personify it as some kind of artificially intelligent, fully autonomous, social actor. Arguing that bots are proxies of their creators addresses this problematic assumption and provides a pathway for future theory and policy focused on addressing the social use of bots.

My informants in this research were clear that they built bots to engage in American politics as they saw it. Their own values, but also the bounds of their software

development skills, were what drove them to use bots for particular means. Walter, an out of work software developer, artist, and amateur bot builder, told me that he was not a great coder but that "new tools like Cheap Bots, Done Quick and Tracery had helped [him] to make functioning social bots" (interview, November 24, 2016). He said he found catharsis in building bots that argued with trolls and particularly outspoken people on the opposite side of the political spectrum from him. He concurrently felt like he was doing something—engaging in activism—by keeping potential trolls engaging with automated accounts and that he was preserving his own emotional well-being by not doing the debating himself. Walter's bots, however, were specifically built to argue with subsets of people with particular views or engaging in conversation using specific hashtags. They were encoded with his values. If I have learned one thing from my research into the use of bots in political communication it is that there is always a person behind the bot.

B. The Democratization of Computational Propaganda

My field work has made it clear that social bots, which allow for more potent coordinated social media activity around particular hashtags and topics, are among the digital tools that loosely organized groups of politically likeminded constituents harness to further their particular political agenda online. These particular bot makers are not powerful political actors attempting to manipulate public opinion on behalf of a government. They are not officially working for election campaigns in bids to drive up social media metrics for a candidate. Many of the people building and launching bots during the 2016 election were members of the voting public.

These users, often acting alone or in small groups, launched armies of political bots made to look like real users but built to post for or against particular candidates or causes. They also launched lone commentary bots. In the U.S., there has been a democratization of who can use bots for political means over social media. This is potentially true of other Western countries, where a combination of regular access computer skills training, high internet penetration rates, and open free speech laws allow people other than elites to make use of computational propaganda. As Arthur, a political bot maker with a day job doing social media strategy, told me, "almost anyone can launch a social media bot" (interview, September 29, 2016)

Research on and with constituent bot builders builds understandings of the democratization of political bot building. It illuminates the ways in which everyday people deploy social bots for political purposes—this includes their discussions on how particular social media infrastructure allows for more widespread use of computational propaganda or automation more generally. The democratization of computational propaganda marks a shift in academic understandings of who uses political bots and how. In contrast to extant research, the political use of automation and algorithms shifts from being primarily used by powerful political actors and towards broader public use.

The democratization of computational propaganda has occurred in the U.S., at least in part, because of the same things that cause the digital divide: access to education, money, infrastructure, and other socio-cultural benefits afforded to people in particular countries. This democratization reinforces research on the digital divide and internet self-efficacy:

"a form of self-evaluation that influences decisions about what behaviors to undertake, the amount of effort and persistence put forth when faced with obstacles, and finally, the mastery of the behavior" (para. 3, Eastin & LaRose, 2000). Those in countries with more access for the internet, and corresponding skills from basic computer use to advanced coding, are more likely to engage in higher level activities—such as political bot building—than are publics in, say, the Global South.

Put another way, the organizational and communicative political power of social media is more likely to be enhanced through social automation by digital constituents in democracies then it is under repressive regimes. In the latter collection of states, past research shows that powerful users, including governments, candidates, and campaigns, are the primary users of computational propaganda (Boshmaf et al., 2011; Woolley, 2016). Siri, a creative professional and amateur bot builder from New England, told me that she saw a shift during the 2016 primaries in the way that politically-oriented bots were being launched online (interview, May 1, 2016). Lots of regular people, she said, were using bots to spread ideas and still others were using them to do other, unexpected, backend communicative activities.

What, though, does this democratization mean for the future of political bot usage? What aspects of social media make this phenomenon, and the broader spread of computational propaganda, possible?

The democratization of propaganda means that constituent builders can, and do, use political bots to launch the disinformation and amplification campaigns now often associated with these tools in attempts to affect political communication. At times, these individuals even set out to purposefully try to alter public opinion related to particular candidates or concepts. However, many constituent builders also use bots in efforts not associated with massively amplified propaganda or attack campaigns. I have spoken to builders who have constructed individual or small groupings of social bots with the goal of making information about political processes more public and easier to understand, of creating art aimed at critiquing particular policies or social issues, and of connecting social or political groups with similar interests but who may otherwise have never communicated.

Roger, a technology sector professional from New York who builds political bots in his spare time, explained to me that it was important how—and for what purposes—bots get launched on social media platforms. He said he could make bots that were as useful in bringing people together over an issue as they were in dividing them (interview, April 18, 2016). Bots, he made clear, are tools that a person uses to convey particular intentions. They are not good or evil, politically motivated or not, in and of themselves. Roger, like many of the people involved in the hobbyist social bot building community, was very wary of calls to ban bots on Twitter outright.

Constituent bot makers like Roger can be divided by two broad, sometime intertwined, motivations for launching political bots and botnets. First, these actors can use bots over

social media as tools for making money. There are now several well-documented cases of users deploying political bots to amplify the spread of disinformation or partisan news in attempts to generate income. Jestin Coler, the person behind the fake news sites the *National Report* and the *Denver Guardian*, has spoken publically about using social media posts to drive users to his sites (Pelly, 2017). For some articles, Coler would garner as many as 8 million views and "with each click, he made money on ads—over $10,000 dollars a month" (Pelly).

Roger told me that bots are often the mechanisms that people like Coler use to get their articles to spread quote. "Bots can be used to plant and boost conversations" about particular topics, he said (interview, April 18, 2016). He told me, and I regularly witnessed evidence to support the idea, that other people use bots to enhance their ability to post messages aligned with their political opinions. That is, they are trying to make a point, rather than to make money. It is true that some are using their profiles to unwittingly market for fake news, but many such users that I spoke to or interacted with believe that they are doing altruistic work to help their country. When asked about the ethics of deploying their own opinions at scale using bots one user, an amateur bot builder from the Illinois, said he was not breaking any of Twitter's terms of service (interview, November 1, 2016). He had, apparently, followed them to the letter. In fact, he said he was using the site just as it was designed to be used because bots have always had a place there.

One hybrid method of deploying computational propaganda which came up in conversations with informants, reports in news stories, and my own observations of bots on social media combines attempts to make money with the goal of seeding and spreading particular political ideas. This tactic makes use of both social automation and human networking. A few months after the 2016 election, I came across information which suggested that small groups of users had deployed covert social cyborg accounts, at times run by bots and at times run by people, to seed political ideas—or links to "fake" news websites—amongst groups of unwitting human users. The human users, in groups thousands strong, would then spread this information in a way that seemed more organic.

Mayer, a computational researcher who claimed to have tested similar techniques, told me that they worked best on Facebook group pages (interview, August 14, 2016). She said this was because those were the areas on that platform where people engage in conversation with people they do not know or do not know well. Moreover, Mayer told me that the cyborg accounts were effective in beginning conversations because people didn't expect to see any automation—let alone more nuanced human-bot communication and corresponding largescale amplification—on Facebook group pages in the same way they did on Twitter.

The engineers of these efforts were apparently able to circumvent algorithms and human oversite that social media companies like Facebook deploy to detect automated behavior and the spread of manufactured trends. This is because they used a degree of human oversight—often logging in through the front-end of the platform—and because these

cyborg accounts only propagated the idea. Once it was picked up by regular users on group pages the cyborg accounts faded out of the conversation. They were often deleted outright.

Organizers of these hybridized attempts to spread political information had success in two areas. First, these tactics allowed political ideas to spread in a way that seemed more natural. There was less chance of accounts or efforts to spread information being flagged by site administrators because the dissemination appeared to be organic. Second, they could mimic grassroots spread of content while also driving regular users to linked sites—in turn increasing advertising revenue for the original architects.

Later, an investigative article covered a similar, though unautomated method, of seeding and spreading content over Twitter. This activity occurred in "pro-Trump 'rooms,' private spaces on conservative Twitter that allow followers to coordinate messages and then retweet each other—dramatically multiplying their impact" (Musgrave, 2017). In these circumstances the goal was to more organically spread provocative political content without, seemingly, making use of bots. Rather "manual tweeting is the cardinal rule" because this better took advantage of the "mechanics of attention" set up by platform algorithms.

Figure 5: Message Spread via Twitter Room

Source: Screenshot from http://politi.co/2viAxZA

Notes: Messages and memes like these often were often crafted on 8chan or 4 chan, cultivated on Reddit, and then spread using either bots or organized private rooms of people over Twitter.

During the course of fieldwork, I manually gathered and saved information on hundreds of social media accounts that used varying degrees of automation to spread information related to the U.S. election. Other research around political bots and the U.S. election have used computational methods to gather information, and conduct analysis on, hundreds of thousands or millions of social media posts (Howard et al., 2016; Bessi & Ferrara, 2016). This latter work is important for several reasons: it generates generalizable understandings of how political bots were used during specific periods of time and establishes broad use of computational propaganda. For this project, however, the goal has been to go deep rather than broad in order to follow particular accounts over time without the additional layer of obfuscation sometimes added in using computation to

massively scale the work. I also wanted to spend time investigating the accounts, making use of manual heuristics for bot detection, in order to connect bots to their makers.

There were several instances in which I was able to deduce the identity of a maker of a particular political bot or botnet. At times, I had luck making contact with builders. One, in particular, stands out. From February to March 2016 I had been tracking several Twitter accounts that seemed to be related. That is, these accounts regularly tweeted highly similar or identical messages, had similar profile and banner images, messaged on similar timelines, and regularly interacted with one another in ways that suggested complicity rather than casual relationships. These accounts would regularly be the first—with seconds of a tweet going live—to like one another's posts, to retweet them, or to comment on them. Each of the accounts overwhelmingly shared articles, rather than generating original content—even in comments. Nearly all of this information was from "alt-right" or far right conservative sources. Most interestingly, and key to the research that lead to locating their owner, was the fact that they all listed the same websites in their bios.

The websites linked by these particular highly automated and highly politicized accounts were all shell pages that linked shoddy advertisements and had no readily visible ownership information. All webpages, however, have searchable information on who owns the site's domain name. This, like an IP address, can be obfuscated through various means. It can also—at times—lead investigators to less careful owners. The sites in question all had the same ownership information. After a bit of searching, I was able to

contact the owner via email. He declined to be interviewed formally, but told me that he was just a regular person with coding knowledge who maintained several accounts on social media for sharing content about conservative politics.

Information about him online revealed that he was a retired software developer living on the east coast. He was not working for a campaign and his accounts didn't share any information in support of a particular candidate, though they did regularly attack Hillary Clinton. After we communicated he listed his real name on the profiles in question—no personally identifiable information had been on the account prior to this. He also made a public point of saying that he not a bot, but a person who believed in conservatism and in upholding the right to free speech. Judging by the rate of posting alone, once every three to five minutes during all hours of the day and night, his accounts used bot software to automate posts and interaction on Twitter. The content of his accounts' tweets and the networks surrounding them also pointed to bot activity. There is a complex relationship underlying the connection between the person using automated software, or bots, on social media and the bot itself. Suffice it to say, this person was a constituent bot user—but he saw his automated accounts as extensions of himself rather than separate entities functioning autonomously.

This was the first of many such interactions. Efforts to connect the makers to political campaigns or organized efforts to produce journalism—false or otherwise—turned up nothing much of the time. There is no reason that everyday people cannot launch political bots on Twitter and it was clear, in many of these circumstances, that this was what was

happening. Conversations with other researchers and journalists revealed that they were having similar experiences in connecting with constituent bot makers. Many were current or former technology sector professionals who were using automation over Twitter to scale their efforts to share political content. One or two even maintained their own servers at home and claimed to have multiple computers set up in order to oversee multiple bot accounts simultaneously. Several, including Mayer, claimed to make use of both automation software and organized human activity to get messages to scale (interview, August 14, 2016).

Some constituent builders turned out to be making single social bots that were oriented towards political critique. Pan, an educator and software engineer living in Chicago, spoke to me about just such a bot that he said he had built as a social commentary outlet (interview, March 8, 2017). This particular bot communicated via Twitter and used automation to scrape information from an open database of facts about nuclear arms and disarmament. The bot would then post particular statistics to Twitter. Pan's hope was that the bot would draw attention to the problem and cause people to think about it more deeply. This particular account has over 50 thousand followers and has been written about in major newsmedia.

Pan calls his interest in bot making an "avocation (hobby) rather than a vocation" (interview, March 8, 2017)" He is a prime example of a constituent bot maker. Though he does not use his bots to spread inflammatory news stories or to make money, he is using his bots to make political provocations. He is trying to communicate about the world

using automation, but his accounts are transparently bots. They do not pretend to be real users supporting a particular candidate or cause. His bots are usually singular, not mutually supportive combinations of interrelated botnets. They work to spread established facts rather than disinformation or defamatory posts. However, Pan is still using his accounts to spread information on a specific cause—much like the bot builders who manage hundreds or thousands of political accounts. His accounts still operate beyond the capacity of a human user. How, then, do his bots differ from amplification or political harassment bots and why does this matter?

First, the fact that Pan builds his bots to be transparent about being bots matters. Bots that attempt to trick human users into believing that they are also human often have the added feature of trying to spread particularly biased of factually inaccurate information. This latter point, about the type of content that the account spreads, is also something that sets Pan's bots apart. His bots access online databases, usually from libraries or museums, and traffic in established facts. Overall, they make political discourse richer rather than interrupting or muddling it. Bots like Pans make it clear that useful bots do exist on social media. Bots are tools for finding, parsing, and spreading information and it is how they are used that really matters in the context of political communication. Calls to ban all bots on sites like Twitter miss the mark because they place the blame on the tool rather than on how it is used.

The bot is a proxy for its creator, not a problematic tool in and of itself—concepts that are unpacked further in the fourth section of this paper. Pan told me that "because bots only

do what people tell them do, they also come with all of the bias of their creator" (interview, March 8, 2017). He noted, however, that "the bias of a bot is hidden by the way people interact with the bot on a site like Twitter." He elaborated, saying that the user sees Twitter bots as mechanical, and thus innocent, functionaries. He said this is not true—the bot is doing what a builder wants it to do. In a sense, though, "the bot gives the coder a bit of distance from the content that shows up on someone's screen. People can't see what is going on behind the scenes, so to speak."

Pan built upon the idea that the bot provides a layer of obfuscation between the builder and other platform users. He said, "if I, or a team of people, was manually tweeting this content other users would expect a quick human response when they asked questions or spoke to the account. They would hold us accountable. It is harder, though, to ascribe intentions to a bot than it is to a human." In other words, even though you can sometimes attribute the bot to a creator, it is several steps removed.

This brings up a number of tricky legal questions surrounding responsibility and intent on platforms where the ability to spread automated political messaging has been democratized. What if the bot, programmed to learn from its surroundings, does or says something unexpected? Also, when it comes to the millions of human users being targeted by bots—which can generate a tremendous amount of noise, can any right to privacy be expected? What happens when this automated noise causes, purposefully or otherwise, democratic activists the inability to communicate or organize?

C. Social media Infrastructure and Constituent Bot Building

What makes a social bot a social bot and why does this matter? What role does anonymity play in the use of these tools? Above all, how do the affordances of social media benefit constituent bot makers? The fact that many bot profiles have little to no personal data on their Twitter account—no last name or no profile picture, for instance—obscures perceptions of legitimacy of such accounts. Some human accounts, however, are similarly anonymous. The anonymity afforded by platforms like Twitter, which lacks a real name policy, and by tools like bots, which act as proxies for human users, makes it difficult to track down the people behind bots. The work that has been done to track the ways that powerful political actors use bots, and to pinpoint just who those actors are, gives researchers clues into how to go about locating the profusion of non-governmental actors—here more specifically referred to as constituents—who use bots in efforts to engage politically online (Ratkiewicz et al., 2011; Metaxas & Mustafaraj, 2012; Woolley, 2016).

Sometimes, as I have explained above, it is as simple as using manual heuristics to track down these particular bot makers. Beyond using names and image searches to match a profile picture with other photos online, links to websites, location data, message content, and follower lists are all salient points of information when working to find out more about the people behind the bots. In fact, when Donald Trump retweeted an account that many users—and news outlets—thought was a political bot account in August 2017, it

was linked to an unofficial Trump merchandise website that led journalists and researchers to the account creators and their backstory (Collins, 2017).

Bots maintain agency in unique ways. Neff and Nagy (2016) argue that their actions can be understood through a symbiotic view of agency. They suggest that this "agency is co-constituted in complex interactions between society and technology" (p. 4917). What society says, does, and expects all have an effect upon what a bot can do—or what a builder wants it to do. Because of the distance between a bot and its builder, and because of the potential for other users' impact upon a bot's inputs and outputs, I argue that the social bot is best viewed as a separate but connected proxy of its creator. This is particularly true when bots are used to enhance the abilities of politically motivated individuals or small groups of citizens because in these circumstances the technological actor is in conversation not just with its builder, but a whole socio-technical system and constrained or molded by the values and beliefs of a particular perception of governance.

Social, and indeed political, bots help builders to reveal, exploit and change aspects of digital systems that otherwise go unquestioned. These changes occur through interactions directed from human to computer, but also those initiated from computer to human. However, it is too simplistic to say that automation on sites like Twitter or Facebook is purely functional. Bots are more than strings of code. The reality of social bot capability exists somewhere in between the poles of bots as wholly autonomous entities or bots as programs or tools. Interviewees have shown a consistent desire to place themselves as

separate but related to, or invested in, the bots they build. They challenge the rigid separation between humans and technologies, makers and mediums

The way a social media site in constructed—the degree of backend accessibility the platform allows public coders, for instance—makes a big difference for how, or whether, bots built by outside developers can function on the site. The way that the social networking occurs on the site also has a large effect on the how successful bots function. The constituent bot makers I spoke to told me that, for these reasons, they most often launched U.S. election oriented political bots on Twitter because of their platform's particular construction and social network. For the same infrastructural reasons, many had similar success in launching political bots on Reddit, Tumblr, and Tinder. Harry, a builder I spoke to early on in my bot research, summarized the reasons for launching bots on specific sites when he said "different platforms, and different communities, have different technical and social benefits for bots" (interview, April 2, 2015). For Harry, Facebook and Instagram make it challenging to launch successful social bots because the former platform's API and algorithms are always changing and the latter does not allow photos to be posted through its API.

That said, several other builders told me about, and were even able to show me, successful Facebook bots. They launched bots through Facebook's graph API or browser emulation, with the latter proving most successful in launching small scale bots that imitated regular profiles. One builder broke uses on Facebook down as messenger bots, bots for data collection, and bots that imitate human users (interview, August 14, 2016).

Some told me that they built to mimic real users in attempts to spread political information. Many told me that they deployed bots to automate political event and group pages. In fact, two different builders told me that when Facebook caught their 'fake' automated profiles the company pushed them to create an automated page rather than deleting the bot outright.

The point about platform architecture making bots more of less feasible is best illustrated by explaining the function of the application programming interface (API). The API is a request and response mechanism for exchanging data with the remote servers where all large online sites store their information. Some companies make APIs with various functions public to external developers in order to make it easier for them to interact with data from the site in question. Researchers, for instance, can access one percent of Twitter data from on a given subject—a representative sample, according to the company—from an API in order to study that information. Accessibility to various APIs can also allow coders to actively launch software programs.

The latter point is underscored by the fact that different ways of interacting with other users on a social platform has an effect on how, or whether, bots can work to drive political communication on the site. On Facebook, users tend to communicate with friends and family or those connected to others one directly knows. It is relatively uncommon to follow or interact with those who one does not know of that platform. On Twitter, conversely, most users interact with a much broader variety of other users, friends and familiar perhaps, but also journalists, politicians, professional athletes, and

yes, random users. The greater flexibility in the way Twitter functions as a social network means that it is easier to launch socially oriented bots that could have success in interacting, and even changing the opinions of, people using the site.

Twitter's decision to have open APIs, in contrast to Facebook's decision to be far more closed in that regard, was originally intended as a democratic function useful to creative users schooled in coding. It meant, and continues to mean, that users can interact with Twitter and the data there in all sorts of ways beyond the manicured frontend pages that most non-coding users see. Developers could launch software in an effort to make the site better for other users or to streamline their own experience. Both Twitter's business site and third party applications make it possible for non-coders to queue tweets over a period of time, rather than having to write from the tweet deck real-time. Access to the API, means that users can enhance abilities like this computationally. They can and do, for instance, write software that automatically creates and posts content—or automatically likes, retweets, or interacts with posts from other users—on its own. This software is a Twitter bot.

Though Twitter's open API access may have been meant to facilitate innovation on the site, it is now the function through which coders use the platform for political manipulation—to spread computational propaganda. Political groups who employ people with sophisticated coding knowledge launch armies of bots to automatically stump for their candidate. Jam, a data scientist who has built bots on to unofficially support specific political parties, compared these political bots to the age-old tradition of political

candidates planting supporters in the crowd (interview, August 18, 2016). Bots, like supporters hiding in a crowd not familiar with a candidate, can make the candidate and his ideas look much more popular than they are in an attempt to get others to join the bang wagon.

As Jam put it, however, bots computationally enhance this ability. Rather than 50 people in a crowd cheering for a policy that might otherwise have been unpopular, you have 5000 bots bolstering every action of a candidate or cause at a rate of interaction far beyond that of a human. Moreover, communication by political bots on social media platforms does not only effect users on the front end by sharing content to convince or making something look for viable then it is through lots of likes or retweets. It also has an untold effect upon the algorithms that sites like Facebook (newsfeed) and Twitter (trends) use to prioritize information on behalf of their users.

Anyone with coding knowledge can launch software using the Twitter APIs. However, it is no longer necessary for users to have a background in a coding language to build and launch simpler versions of the same software that drives the most sophisticated political bot armies. Walter, another bot builder—or perhaps, in this case, bot user, told me that he had never learned to code (interview, October 24, 2016). He uses to third party platforms to launch his Twitter bots. Free online services such as "If This Then That" (IFTT) allow those with little to no coding experience to create applets, which use conditional statements to computationally interact with sites from Gmail to Reddit. More explicitly

political groups of constituents, such as Patriot Journalist Network, provide "members powerful tools to facilitate more effective and influential use of Twitter" (PJNet, n.d.).

In contrast to Twitter, Facebook requires more coding knowledge, and a more innovative approach, to launch successful political bots. Bots are also less scalable on Facebook. Mayer told me that on Twitter she could run several thousand bot accounts, while on Facebook she was lucky if she could maintain 15 (interview, August 14, 2016). When her bots were caught, as mentioned above, Facebook suggested they be converted to pages—where automation is allowed to a greater degree. Interestingly, these are likely the areas where bots would likely have the most effect in spreading biased ideas or misinformation on Facebook because they are where people have conversations with people—even strangers—outside their close social network.

D. Bots as Proxies for Human Builders

In addition to spending nine months talking with the people building and using bots during the 2016 U.S. Presidential Election, I have another spent four years interviewing bot makers around the world. This work has been done in conjunction with two connected projects at the University of Washington and the University of Oxford, concerned with the study of what our team has termed "computational propaganda" (Woolley & Howard, 2016). The goal has been to develop understandings of social automation and political communication online. In some ways, this work has been a theoretical endeavor aimed at understanding the particular capacities of social bots as social actors. Bots can be used as simple tools on social media platforms, to automate

pre-programmed messages on Twitter or to moderate posts on Reddit. They can also, however, be drivers of sociality in unique ways. They can be constructed to learn communication skills from other users, an unpredictable venture in and of itself, or be deployed and left to interact with their environment with little oversight.

The empirical goal of my own research on the 2016 United States Election has been to understand these unique automated tools by talking to the people who build them for political purposes. I interacted with makers building and deploying bots over social media for all sorts of reasons: to curate news, to manipulate political conversations, to create political art, to comment on social problems, and more. These interactions, and the projects in question, have revealed a special relationship between bots and the people who build them. Bots act as a proxy for their makers; they help them to creatively interact with digital systems in surprising, and often surprising, ways. They are agents of their builders, but they are also separate.

There is a tendency to personify the social bot. Social media users talk about creations like @OliviaTaters or Xiaoice as if they were thinking automatons, human-like artificial intelligence akin to C3PO or eccentric robot helpers more like R2D2 (Lloyd, 2016). Creators often give their bots human names and creative personalities, they speak about them like friends or pets. Social bots, however, are still built to do tasks specified by developers and even advanced machine-learning bots, such as Microsoft's @Tay, reveal the imperfections and biases of their builders (Neff & Nagy, 2016). Most bots on social media are semi-automated tools that allow the person who launches them to interact with

platforms in both an obscured infrastructural and shared interactive capacity. In other words, social bots communicate with servers, applications, and databases on the 'backend' of websites, but they also interact with human users on the frontend of those same sites.

Thinking of bots as proxies for their creators means that they act as surrogates, but also that they function as entities that enact change on a system while they themselves are also changed. Bots that interact with complex social systems have the ability to do things beyond the expectations of their builders, to achieve the unexpected.

Harry was one of the first builders to give me an explanation of the complex relationship between the bot and its builder (interview, April 2, 2015). Referencing the work of Ian Bogost, he stated, "bots are independent objects that do the work of the philosopher" (Bogost, 2012). This is a concept that requires some unpacking. First, Harry wanted to make it clear that bots are separate from their creators. They are autonomously functioning. Second, he elaborated on the position of the philosopher-builder as someone who builds functions into these digital objects and then relinquishes the majority of control over their activities in the world. Yes, bots do the work of the person that constructs and releases them, but they also have their own identity, they interact with the Internet in ways removed from that builder.

Bots role in pushing human social goals is something that should be complicated rather than simplified. It is not only a tool, but it also is not sentient or emotive. Bots mirror

the builder, while also mirroring the networked computational systems in which they operate. After this, and many other, discussions with those who build and launch social bots on social media platforms we have come to see bots as extensions of their creators. A bot stands in, and often does the will of the person who built it. However, a bot can also challenge the authority of its maker. It functions and interacts with its environment, independently and automatically. The builder is not the bot and the bot is not the builder. Rather, the bot is a separately functioning representative of its creator; the bot is a proxy.

In the course of my interviews with bot builders, and in my work revealing the manipulative political uses of bots during elections and security crises, I have regularly encountered conversations about the complicated user-bot relationship. During the 2016 Presidential Election collaborated on work with the team at Oxford to identify certain automated Twitter accounts working to amplify partisan points of view in attempts to manufacture trends and project consensus (Howard, Kollanyi, & Woolley, 2016). However, I received pushback from users behind accounts reported upon as being bots or highly automated.

One user, a developer from the United States, wrote an email to me in which he stressed that he was not a bot, but a person who used software to automate his presence on Twitter. This particular user had, in his relatively brief tenure on the platform, produced hundreds of thousands of messages. Most were heavily in favor of a particular political perspective or, at times, heavily against oppositional views. The account tweeted nearly every second, with a high degree of topical specificity, both markers of automated

activity. It was, however, diverse in the way it interacted with content and other users: it was able to both like and retweet existing posts but also to generate original messages.

It would be true to say that this user was harnessing a bot in order to bolster his online presence. By definition, a bot is a piece of automated software used to do automated tasks that a human would otherwise have to do. But interaction with him complicates a mutually exclusive division between him and the software automating his profile. He saw his profile as being reflective of his views. For him the software was mostly functional, it allowed him to have some kind of online presence at all hours. The software, however, is not the person. Or, put differently, in this case it was the software, and not this user, that was actually generating tweets and interacting with site. But, the bot didn't exist by itself. It was made as an extension of its creator. While people shape the content and trajectory of automation used to streamline interaction on and over social media, they are also separate from that technology and vice versa.

Anecdotes like this reveal the usefulness of thinking of bots as proxies. The word proxy describes organizations, activities, and people who "act as a substitute for another" ("Proxy", 2016). I use this definition when referring to a social bot as a proxy for a human user. The bot is endowed with agency or power, and can exercise that agency on behalf of the builder. Because a person must construct the code that allows the bot to function, they can have a degree of faith—and culpability—in what the bot does or will do. However, the expectations of creators are challenged on social media platforms where the bot interacts with other people or programs. Here a diverse range of inputs and

outputs can affect the actions of a bot. In short, is it hard to predict exactly how any social program will function under different stresses. Tay was an example of this. The case of a bot built by developer Jeff van der Gloot (Weaver, 2015) wherein the bot said it wanted to kill people, stands as another. Here, a proxy behaved in a way that its creator didn't intend. The bot went beyond the tasks its builder believed he had programmed.

A social bot is, therefore, an automated software program built to stand in for, represent, or extend a human user. These digital automatons allow users the unique ability to interact with both systems and social spheres in ways that they otherwise couldn't. Bots allow users the ability to act exhaustively. While a user has to sleep, a bot does not. They allow users to multiply their presence online, as one person can control many bots. They allow users to speed up their actions online, to computationally enhance the ability to complete tasks. They allow users the ability to veil their actions, providing a degree of separation and anonymity between the user, the message sender, and their target, the message receiver. They allow and facilitate communication; they mediate conversations and can be seen as a new medium in their own right (Woolley & Guilbeault, 2017)

Bots are infrastructural elements of social media systems built on top of existing platform. The work of deploying a bot over Twitter or Facebook does not require direct collaboration or involvement of platform designers or engineers in order to change or allow for original features that users want to make possible. For example, users can build and deploy bots to translate text from one language into another. This accessibility can be harnessed in very practical ways. Several interviewees argued that they were most

generally using bots to add additional functionality on systems already in use. Still others argued that this ability allowed for genuine creativity. Bots have been used to generate media-oriented art based upon user's profiles and preferences, to critique news media, even to argue with other users. Both varieties of bots offer new possibilities for interaction on, and with, social networking sites.

When a person types and sends a tweet, it is clear that they're motivated by a desire to convey a thought or feeling. But, when the text is produced through combination and sent automatically by a computer script, these motivations become less clear. This means that it is exceedingly important to look beyond the content of automated messages to the bot designers, themselves. While it is possible to interpret the goals of a bot based on its messages, talking to bot makers helps to understand what their goals—elements that have a meaningful impact on the form and features of bots. Through interviews with bot makers, I discovered that bots both enable and extend the communicative reach of their builders. Interviewees made it clear that social bots function well beyond the purview of simple automated messaging, that they have a unique capacity for complex interaction. Because some bots are built to learn from their ecosystem, they often display startling characteristics borne from factors outside of their original code. They are not human, but they are invested with power that can be directed diversely and unpredictably.

The twitter feed chronologically displays content from followed users, and retweets and @-replies from those users. If bots can garner followers, they can count on their content being in feeds. They can also count on their information to be consumed and repurposed

by other users. If they are built to learn from their surroundings they can be programmed to draw behavior from their followers. Otherwise, they have two options. 1) using popular channels or topics to connect with users or 2) speaking loudly to get attention.

Examples of the first category of bots includes listener bots, which wait for a particular topic to come up and then center their communication attempts upon users discussing that topic. Pan noted that one way to get people to pay attention to a bot is to have the bot listen for users speaking about politics and then engage those users (interview, March 8, 2017). This is especially useful if the builder is working to critique ideology, policy or politicians. Makers have also used these bots to bridge the gap between communities that might not otherwise speak or connecting isolated users to information to which they might not otherwise have access. These bots can leverage databases of information while also monitoring newly created information online. They then share that information to make a clear political statement or draw attention to an action or issue through a proxy relationship to their builder.

Other listener bots can monitor sites and databases for manipulation or wrongdoing. Wikiedits bots, for instance, are often created with to track politicians' edits to Wikipedia pages (Ford, Dubois, & Puschmann, 2016). Roadblock bots, another type of listeners, are built to generate noise. They monitor feeds for a particular topic and then attempt push users from away from it in order to distract them or prevent them from organizing online. Such bots have been used to demobilize activists trying to organize and communicate on Twitter but can also be used to drive traffic from one cause, product, or idea to another

(Woolley, 2016). These bots can do more than drown out activists or drive people to ideas. They can also be used to direct targeted noise at users who attempt to harass other users.

The bots in the first category, bots that listen and then engage, can be bridged to the second kind of bots, those that work to garner attention by being loud or provocative, by a type of bot that one interviewee called "dancing bare-ware." This variety of bot is built to draw attention by saying strange, quasi-human, phrases. Like a dancing bear they engage in unexpectedly human behavior in order to get other users to engage with them. Dancing bear-ware can, however, be built to listen and then engage on particular topics. Such bots have been used to argue with people about politics or distract those who engage in racist behavior.

The second type of attention oriented bots shout at someone until they listen. They can speak directly to a user through @-ing them or tagging them in images. This means content shows up in mentions, and is potentially displayed to others in their network. The number of people involved in this tactic used to be limited by the twitter character limits, but as this chapter was being written @-replies have been excepted from character counts leading to potentially infinite number of @s. An example of this type of bot was given to me by a respondent, who said they had tracked a bot-driven smear campaign. In this instance, a builder or group of builders deployed hundreds of bots to tweet rumors about another user. These tweets were feasibly visible to anyone who looked the embattled users profile. Though distressful, the attack was ultimately not very successful in

mimicking a crowd of real human attackers because the bots were newly created and had no followers, and thus no audience. “In order for bots to be effective,” Walter told me, “you have to be able to verify you message is reaching people” (interview, October 24, 2016).

E. Digital Constituents and the Diffusion of Political Bot Usage

In this chapter I have described the ways that bots get used by digital constituents. This actor group is best described as a broad grouping of individuals in the U.S. who are engaging in the use of political bots but who are not directly connected to campaigns or media institutions. These political bot users are also not, generally speaking, tied to NGOs or other civil society groups. They are regular people who make use of political bots in attempts to get their political views across. At times, they use political bots to make money by spreading links to fake news—or other junk content—websites populated, in part, by advertisements. Digital constituents use political bots in attempts to manipulate public opinion and spread disinformation, but they also use them to spread genuinely useful political information and to critique power.

I argue that the use of political bots by digital constituents has led to a democratization of propaganda. Past research argues that political bot usage is the pursuit of powerful political actors: governments, militaries, politicians and their campaigns, political consultants, or political parties. My data reveals, however, that everyday citizens in the U.S.—people from many professional backgrounds and parts of the country—use political bots. The same social properties that contribute to the digital divide—greater

wealth, more reliable access to education and other social resources, and better internet connectivity in developed countries—leads to the democratic usage of political bots in the United States by particular groups of people. Digital constituent political bot users are most often software developers or work in other professions related to computer coding, social media, or the internet. Some of these actors do not code, but use third party tools to launch political bots.

I also suggest that social bots—and political bots more specifically—are proxies of their creators. Bots are not independently functioning, artificially intelligent, agents. They are intrinsically connected to the design, ideas, and will of their builder. They are deployed to function automatically and they can sometimes interact with their environment in surprising ways. However, there is always a person—and a person's intent—behind a bot.

Chapter 6

Conclusion:

Disinformation, Anonymity, and Automation

A. A Familiar Story, at Scale

People that use social media such as Facebook and Twitter are familiar with this character: the person who always seems to be using these sites to argue about politics, no matter what the context. Many scholars have discussed the democratizing potential of the internet. It allows regular citizens to become documentarians, podcast hosts and investigative reporters. It provides the public in countries with open internet access, including the U.S., the means to explore previously unthinkable quantities of information. It also allows that political junkie acquaintance the ability to become an amateur political pundit over social media.

In some ways the use of bots, especially as seemingly captivated and engaged groups of followers, allows for this punditry at scale. On sites such as Twitter, which have open APIs and flexible automation policies, this has led to democratized use of computational propaganda. During the 2016 U.S. presidential election people across the country, not just power political actors, made use of bots to spread political content ranging from fact based journalism to outright disinformation. I spoke to one man, a citizen bot maker living in the Midwest, who maintained several servers and a fleet of computers in his own house. He was just this kind of citizen propagandist.

Computational propaganda is still propaganda. What has changed is about this new form of an old strategy of control and coercion is that it now happens at scale. This new political communication strategy is scaled not just in terms of the number of people who places around the world that can and do access it but also by the computational power and constantly advancing software—including bots but also data analytics technology including machine learning and sentiment analysis—that facilitate it.

Two primary features of computational propaganda separate it from traditional propaganda: anonymity and automation. Automation allows all sorts of politically inclined people to amplify their views. It also allows them, by correlation, to suppress the ideas of others. This is especially useful for well-resourced political actors—candidates, campaigns, and political interest groups—who can buy custom built bot armies to support their agenda or attack their opponents. Anonymity is perhaps even more revolutionary in its capacity to allow for the profusion and success of propaganda. Not only can bot builders scale their work, but they can do so without fear of being found out. The same feature of certain social media sites that allows democratic activists the ability to mask their identity in oppressive regimes, is a critical tool for those in democracies hoping to manipulate public opinion.

Olaf, the Midwestern bot builder, used both anonymity and automation to great effect in his efforts to spread his political ideas (interview, September 17, 2016). Imagine that person on Twitter who is always talking about politics, now imagine them with 2,000 accounts set to do their bidding—that's Olaf. Olaf claimed that he ran several thousand

different Twitter political accounts using automated software. He showed me several of these accounts—all had tweeted hundreds of thousands of times in their short lifespans and some over a million. He told me that his goal was to use Twitter to become a pundit on traditional media. He thought that if he became famous for his political opinions on Twitter than he could become famous on TV.

Olaf was unique in his ability, especially as a digital constituent, to use bots to massively magnify his own political communication online. Most other digital constituent bot builders I spoke to ran ten, twenty or maybe one or two hundred accounts on Twitter using bot code. Olaf, though, was a retired software engineer. Candidates, campaigns, and media companies—however—have no issue maintaining thousands of accounts set to do their bidding. Militaries, including the U.S. armed forces, have pioneered the use of fake profiles—or persona management software—in information warfare (Rockefeller, 2011). Security organizations and sovereign states worldwide have followed suit, and many use and maintain botnets in their efforts to protect and enforce security in the ways each sees fit (Bradshaw & Howard, 2017).

Though the majority of the U.S. public—and certainly U.S. policy makers and technology leaders—are now well aware of the problems posed by "fake news" and bots, many of these problems lack current solutions. Twitter and Facebook have cracked down on the usage of bots and the spread of disinformation, but neither these techs giants nor their users alone can prevent the spread of computational propaganda. Innovation in this space continues, with the rise of so-called "deep fakes" during elections and crises—

doctored images and video with the intent of manipulating public opinion—and even disinformation via virtual reality (Gupta et al., 2013; Wardle, 2017). Moreover, as users move to close system chat apps such as Telegram and Signal—which many presume to be more secure than their more open counterparts from Instagram to Reddit—so too does disinformation migrate to these platforms (Chesney et al., 2018).

B. Confusion over Coercion

In this research on the people who build and used bots during the 2016 U.S. presidential election I often came across the perception that political bots are generally used to coerce people into changing their opinions about political ideas or candidates. Many people I spoke to, even experts in computer science, data journalism, and political campaigning, made the argument that political bots are mostly used over social media to artificially amplify partisan news, disinformation or the followings of particular profiles on sites like Twitter. While this is true, the bot makers building the most advanced automated products—those that learn from their environments and respond accordingly and those oriented towards platform specific trends rather than bombing thousands of profiles with information—asserted that much of their work was oriented towards confusing the democratic process rather than coercing citizens (interview, August 18, 2016; interview, November, 11 2017).

Advanced bot makers, those thinking about the future of computational propaganda, told me that they believed that confusion—though a less direct route to allowing success of their ideology—was a potent means for control (Ziggy, 2016). Confusion, they told me,

allowed for a vacuum in communication that could be filled by groups or candidates speaking the loudest (whether through automation or bombast). They drew upon the ideas on disinformation as a mechanism of confusion and control pioneered by the Russian government—and practiced by that state in an attempt to manipulate communication during the 2016 U.S. election (*Washington Post* Staff, 2017). Confusion will likely play a critical role in future attempts to promulgate computational propaganda.

Tactically, those seeking to coerce, confuse and, ultimately, control, spoke about the fact that political bots did not need to be built to mimic real people or communicate with real people (interview, November 1, 2017). They were also building bots on sites like Twitter as shell accounts—lacking profile pictures, about sections or robust followings—set to communicate with the backend code of social media sites rather than with frontend users. Builders spoke about botnets constructed to constantly tweet using specific hashtags in attempts to get those hashtags, and associated subjects, to trend. The logic was not only that more users on social media might see the trend because it had been reprioritized—curated—by the social media firm's algorithm, but also that reporters would pick up the trend and produce content related to it. Indirect methods of computational propaganda, or using code to communicate with code to communicate with people rather than using code to talk to people, are on the rise.

C. Methodological Constraints

This study uses fieldwork amongst bot builders and trackers—namely interviews, participant observation, and deep hanging out—to build understandings of the usage of

bots and computational propaganda during the 2016 U.S. presidential election. These methods were directed towards studying how three particular actor groups used these emergent and evolving tools of political communication: candidates/campaigns, journalists, and digital constituents. During conversation and interaction with the people surrounding the campaign, and with knowledge of these particular tools, I also worked to gather data on the various bot accounts that these builders showed me. This ethnography of information, or technology, explored the ways the technical creations of these developers both reflected their builders but also drove sociality in their own unique ways.

This project is fundamentally qualitative. Where a quantitative study might reveal the breadth of bot usage on social during a particular event or surrounding a particular subject, this research is focused upon building in-depth understandings of the communicative and social processes surrounding the use of bots and computational propaganda. This work is, therefore, not reproducible or generalizable—though it is informed by empirical ethnographic research. It is descriptive and speaks to what happened surrounding the use of bots and computational propaganda during the 2016 U.S. presidential election but it does not make causal inferences. It is useful in making sense of the people who built and used political bots during the contest, and what they intended, but cannot create a generalizable formula that supports measureable effects upon broader populations.

Other scholars concerned with the political use of bots consistently ask how bots and computational propaganda can be used to effect people or the general spread of

information online(Ratkiewicz et al., 2011; Bessi & Ferrara, 2016b; Howard et al., 2016). This study is not concerned with these types of effect. It is focused upon studying the mechanisms by which computational propaganda, and political control via bots, are produced. The goal from the outset has been to provide research and analysis that builds understandings of a new societal and communicative phenomenon via the perspective of those who have built the technology that drives it, rather than to establish the effect computational propaganda or bots have upon political processes such as voter turnout or voter behavior.

This finding of this study are limited by the fact that they are derived from information from three actor groups. This study is not useful, for instance, in identifying who builds and uses bots and computational propaganda and why on behalf of corporations or NGOs. Future studies might focus upon builders who work on behalf of these group, and indeed some researchers have already found evidence of corporate and special interest bot usage on (Metaxas & Mustafaraj, 2012b; Gorwa, 2017).

The myriad positive democratic uses of bots, while discussed in the chapters on journalists and digital constituents, are not the focus of this paper. In some ways, these positive uses were overshadowed by the fact that the research for this particular project was focused upon a particular contentious election. Research has demonstrated that bots can be used as social scaffolding for online groups who might not otherwise connect, and further work has been done on additional benefits of journalism bot use (Hwang, Pearce, & Nanis, 2012; Lokot & Diakopoulos, 2016; Woolley et al., 2016). The democratically

beneficial uses of bots might be better explored in a comparative analysis of cases particularly focused upon this technological capacity.

D. Reflexivity and Epistemological Considerations

The epistemological framing of this study, which prioritizes a critical interpretive framework over one more post-positive, might also be considered limitations in the breadth of applicability to more general online communication around the globe. I approach the study of communication, and indeed of bots and computational propaganda, via an interpretive framework of knowledge that prioritizes empirical observation and participation while maintaining a critical emphasis. This focus is the result of a mixed disciplinary background, which combines both traditional and radical approaches to the study of anthropology, cultural studies, and communication. While studying within each of these fields I have undertaken what Davis & Jasinski (1993) call culturally situated communication research.

From anthropology, I have drawn on empirical participant observation, conducting research in the natural environment in order to understand contextual face-to-face interactions. From cultural studies, I have foregrounded the political nature of culture in this research alongside the desire to interrogate the site of both production and consumption of bots and computational propaganda during the 2016 election. This political commitment extends further in that it has led me to seek out this research project because it, in many ways, addresses social concerns and questions of inequality. Research on bots and computational propaganda have provided me with what Schramm (1983)

calls a ‘problem,’ a problem of communication. This problem is academic rather than literal and inspires me towards the acceptance of transdisciplinary and a multi-faceted approach to understanding interaction and the sending/receiving of messages.

Because of these meta-theoretical commitments I am tied to neither purely post-positive nor purely interpretive/critical concerns, though I am skeptical of truly positivist research because I view knowledge as situational and value-laden. I, like Craig (1999), believe many traditions of research can be used to answer questions about communication processes. This is true when it comes to traditional media and also when it comes to new forms, including social networks and bots. Like Bennett & Iyengar (2008), I believe that the combination of technical innovation and social circumstance can alter the way we understand the potential effects of certain types of communication. Theoretical shifts like this have led communication researchers away from concept of strong media effects and back towards a minimal effects model, suggesting that social theory is both versatile and volatile (Bennett and Iyengar, pp. 707-709).

Because my research is interdisciplinary, and my epistemology interpretive with a leaning towards empirical data gathering, I have been afforded the opportunity to work with scholars from various backgrounds. These interactions continually help me in understanding diverse points of view but also provide challenges which have, thus far, been mostly positive. At times my meta-theoretical commitments lead to debate with scholars who lean towards positivism, but I tend to find these conversations more constructive and enlightening than perplexing.

This particular project has allowed me to work with academics in the computer, information, and (various) social sciences, critical scholars, legal experts, journalists, people who represent civil society organizations, and those involved in political and policy oriented groups. These interactions allow me to develop a broader understanding of bots and those who make them, particularly those who use them for political means.

While transdisciplinary work outside of a communication department could provide challenges to future research, it also aids this project in that it provides understandings of problems associated with interactions between humans and machines. Theorizing about social role of intelligent technology, and situating this understanding within communication research, is challenging. I believe, however, that communication is a fitting venue within which to study such a phenomenon because, as many scholars above note, it is a fundamentally diverse field.

E. Frontiers for Political Bot Usage

Several areas are at the forefront of innovation, and future problems, associated with political bot and computational propaganda usage. The first, both in the U.S. and globally, is policy and the law. How will these political domains be affected by the rise in political manipulation over social media? What laws are needed to regulate firms where disinformation is spread? The academy must aid policy makers by undertaking more empirical research to inform policy recommendations to be delivered to key U.S.

politicians, policy experts, civil society groups, and journalists. Campaign finance, election law, voting rights, privacy, and several other areas of the law are currently being affected in both unforeseen and complex ways by the spread of political disinformation over social media (Howard, Woolley, & Calo, 2018). Solid research into the ways computational propaganda contravenes the law is a crucial step in addressing the policy gap at the intersection of information dissemination, automation, social media and politics.

We need better software, informed by both social and computer science research, to help researchers, journalists and activists keep up with the challenges posed by the modern disinformation threat. Tools could include high-powered data intelligence platforms, that make use of bots in parsing large sets of relevant data, usable by these groups worldwide. They ought to exploit recent advances in graph databases and machine learning, and cheap, massive computation to dramatically accelerate investigations. The target should be to accelerate civil society in identifying patterns of activity that would help to root out entities backing disinformation campaigns, in addition to uncovering a great deal about when and where these campaigns are occurring.

Longitudinal work can help establish more solid metrics for tracking information flows related to this the use of political bots, computational propaganda, and, correspondingly, disinformation and online polarization. Quantitative insight into the roles of automation, network structure, temporal markers, and message semantics over social media can allow experienced researchers to effectively create ways of measuring the flow of political

manipulation over social media over sustained periods. The results of longitudinal research on this phenomenon will be crucial to building evolving long-term public and governmental understandings of computational propaganda.

F. Studying Bots as Communication and within Communication

Bots, in the simplest sense, are tools for communication. Political bots, which mimic real people on social media, can be considered a new form of media because they are used to transmit information from their builders to broader communities. Bots, and the complex ways they facilitate political communication and challenge notions of communicative agency, are only just beginning to be explored in terms of communicative utility. The discipline of Communication must interrogate both bots generally and political bots specifically as both emergent, important, tools and unique semi-human actors. This work should be informed by extant work from the discipline that explores questions related to communication, agency and the digital sphere.

Gillespie's (2011; 2014) work on social media algorithms will be useful in thinking through how political bots' automated communication on Twitter—much like a trending algorithm's automated decision making on the same platform—can obscure political intentions. Both bots and algorithms are used to spread information in a way that makes human labor behind that communication obscure. It is crucial that communication scholars consider the fact that bots are encoded with politics by the people who build them. They are not apolitical entities that simply transmit content. That said, like algorithms, they are also manipulated by people beyond their builders. Other social media

users can game both algorithms and bots in order to achieve particular ends. The case of Microsoft's Twitter bot Tay, which unexpectedly repeated racist content from other users, highlights how online communities can get chat bots to do and say things unintended by the creator.

Vertesi's (2008) work is useful in drawing attention to the fact that bots can change perceptions the information they are intended to spread. She has done significant work around the socially imbued concept of techno-scientific representation. She argues that formal and practical representations of objects (for instance, the London underground map) often diverge from the colloquial and enculturated view of said objects. In other words, technological mediation can change the way people experience things in their everyday life but so too can people remix the 'planned' meaning of producers. Bots, like Tay, are often build with a specific use in mind but are often gamed by the online communities in which they are launched.

Vertesi (2015) continues the work of speaking to critical issues of representation in her ethnographic work with the Mars Rover team. Not only in this work a multi-site and multi-method endeavor of capturing a complex social arrangement of professionals from different backgrounds, it is also one of interrogating the interactive role of scientific visuals during the Rover mission. Vertesi places importance on the analytic framework of 'drawing as,' in which scientists use digital technologies to draw the red planet in ways that have implication for future research and viewings (22). Vertesi argues that 'drawn' images of Mars must adhere to communal norms and 'constraints' in ways that shape

both the subjects and objects of analysis (pp.191-193). This framing is crucially useful in theorizing about the relationships of bots and their builders. How might future ethnographers work to illuminate the normative values of the bot building community? This work could, in turn, promote new uses of automated communication technology.

G. Defining the Bot

At the beginning of this research I thought of bots using a fairly broad categorization and of political bots much more narrowly. Bots, as I defined them for my research, were nearly any and all automated programs used to digital tasks online. Political bots were more specifically those automated accounts used on social media to mimic real people in attempts to manipulate public opinion. My work with bot builders—whether campaign builders, journalist builders or digital constituent builders—has helped me to understand both that these different groups define bots different but also that scholars of communication need to consider more narrow definitions of different types of bots in order to be successful in discretely analyzing their usage in particular communicative contexts.

With this in mind, it was much more useful for me to consider how different actor groups used political bots rather than how these same groups used all bots for political communication. In other words, the narrow definition of the political bot allowed me to more effectively interrogate a particular type of emergent political communication over social media then did the broad definition of bot. In the beginning of this research many of my informants challenged me on what I meant when I said "bot". As time went by it

became clear to me that my own diverse use of the term was informed by a similarly diverse use of it by other scholars and within popular culture. With this in mind, I would now argue that communication scholars must clearly define the bots that they study in terms of their particular function in spreading information.

For this research, I was originally most interested in how bot builders used their bots to generally effect political communication during the 2016 contest. After I finished my data collection and analysis I realized that the data revealed more interesting findings related to how particular actor groups—from campaign consultants to digital constituents—built bots to mimic people. With this in mind, I plan to refine the way I scope definitions of the bots I study in future research. The definition of the political bot proved useful in this research, the definition of bot proved confusing.

H. Challenges for Ethnography of Information

Spending time with non-human or semi-human actors is a challenge for ethnographers of information—who hope to study what happens when technological actors are driving sociality in unique ways beyond those envisioned by their builders. In this research, I worked to explore how an ethnographer might do deep hanging out in a community at least partially composed of political bots. I followed particular accounts over long periods of time, gathered data from these accounts, and took notes on how they engaged in political communication. I also, however, spoke to the people who made and built these bots. In some senses, then, I practiced both traditional ethnography amongst humans engaged in communication alongside ethnography of information.

At the nexus of these two field oriented methods is a challenge for scholars who hope to illuminate the communication practices of bots or, for that matter, those of algorithms, artificially intelligent robots, or information technology systems more generally. Specifically, how is ethnography of information both different from and similar to traditional ethnography? What can researchers learn from spending time with both makers of tools and the tools themselves? Beyond this, how might we think about bots as something other than simply tools—especially when they are driving sociality in ways unintended by their creators?

In this work, I found myself falling back upon exploring the difference between what political bots did online during the 2016 contest versus what builders intended these bots to do. More often than not, however, I relied more upon data from human builders than I did from data from bot accounts. This was potentially because I am better versed in traditional ethnography than I am in the emergent practice of ethnography of information. It was also because I found it overwhelming, and at times futile, to qualitatively track and interact with the tens of thousands of political bot accounts that had unclear origins, were consistently being deleted and communicated in seemingly frantic non-sequiturs.

The ethnographic study of political bots on sites like Twitter, Facebook and Youtube is a challenge for the very reason that bots are not human and ethnography is built around the study of human culture, of the norms, values and beliefs of groups of people. But bots are often build to mimic people, and they certainly communicate with people. I've

consistently found that they do things their builders don't plan for them to do. Other people online find ways to get social media bots to do things never expected by the wide variety of groups who build and launch them. But bots also, regardless of human based inputs and outputs and because of other unforeseen technological communications, do things that are unhuman that still drive sociality of both human and non-human communities.

Therein lies the challenge. How can scholars of communication and other social sciences effectively study bots using ethnographic methods when they don't know who built the bots in question? Comparison between the builders and their creations in the wild is not always—in fact, rarely—possible and this means we need other ways to do ethnography of information. We must ask ourselves: what is interesting about the ways that bots and other information technology systems communicate and socialize outside of those actions intended by human builders? Studying builders is a great entry point into understanding the use of new, or newly important, communication technology. But beyond this lies the more pure analysis of the automated systems themselves—apart from people or in conversation with people who had no hand in their creation.

References

AJ. (2015, March 31). Interview D.C. 1.

Al. (2016, October 12). Interview New York 8.

Albright, J. (2016, November 11). What's Missing From The Trump Election Equation? Let's Start With Military-Grade PsyOps. Retrieved May 22, 2017, from https://medium.com/@d1gi/whats-missing-from-the-trump-election-equation-let-s-start-with-military-grade-psyops-fa22090c8c17

Allcott, H., & Gentzkow, M. (2017). *Social Media and Fake News in the 2016 Election* (Working Paper No. 23089). National Bureau of Economic Research. Retrieved from http://www.nber.org/papers/w23089

Anderson, K. (2016, July 11). Microsoft CEO Satya Nadella says chatbots will revolutionize computing. Retrieved January 30, 2018, from https://www.onmsft.com/news/microsoft-ceo-satya-nadella-says-chatbots-will-revolutionize-computing

Andrews, N. (2016, February 25). Pro-Trump Twitter Bots at Center of Nevada Mystery. *Wall Street Journal*. Retrieved from http://blogs.wsj.com/washwire/2016/02/25/pro-trump-twitter-bots-at-center-of-nevada-mystery/

Andy. (2015, March 15). Interview Budapest 2.

Angelo. (2017, May 4). Interview Online 5.

Arnaudo, D. (6/17). *Computational Propaganda in Brazil: Social Bots During Elections* (Computational Propaganda Working Paper Series No. 2017.8) (p. 39). Oxford, UK: Oxford Internet Institute, University of Oxford.

Arthur. (2017, September 29). Interview Online 2.

Associated Press. (2016, March 3). Transcript of Mitt Romney's Speech on Donald Trump. *The New York Times*. Retrieved from https://www.nytimes.com/2016/03/04/us/politics/mitt-romney-speech.html

Barot, T. (2016, December). The botification of news. Retrieved November 7, 2017, from http://www.niemanlab.org/2015/12/the-botification-of-news/

Beckett, C. (2016). What does the Trump triumph mean for journalism, politics and social media. Retrieved from blogs.lse.ac.uk/polis/

Benkler, Y. (2006). *The Wealth of Networks: How Social Production Transforms Markets and Freedom*. New Haven, CT: Yale University Press.

Bennett, W. L., & Iyengar, S. (2008). A New Era of Minimal Effects? The Changing Foundations of Political Communication. *Journal of Communication*, *58*(4), 707–731. https://doi.org/10.1111/j.1460-2466.2008.00410.x

Berry. (2016, March 8). Interview Michigan Two.

Bessi, A., & Ferrara, E. (2016a). Social bots distort the 2016 U.S. Presidential election online discussion. *First Monday*, *21*(11). https://doi.org/10.5210/fm.v21i11.7090

Bessi, A., & Ferrara, E. (2016b). Social bots distort the 2016 US Presidential election online discussion. *First Monday*, *21*(11). Retrieved from http://firstmonday.org/ojs/index.php/fm/article/view/7090

Bialik, C. (2016, April 14). Everyone Has Fake Twitter Followers, But Trump Has The Most. Sad! Retrieved July 24, 2017, from

https://fivethirtyeight.com/features/everyone-has-fake-twitter-followers-but-trump-has-the-most-sad/
Bilton, N. (2016, August 5). Trumps Biggest List? The Size of His Twitter Following. *Vanity Fair*.
Bogost, I. (2012). *Alien phenomenology*. Minneapolis, MN: University of Minnesota Press.
Boshmaf, Y., Muslukhov, I., Beznosov, K., & Ripeanu, M. (2011). The socialbot network: when bots socialize for fame and money. In *Proceedings of the 27th Annual Computer Security Applications Conference* (pp. 93–102). ACM. Retrieved from http://dl.acm.org/citation.cfm?id=2076746
Bowler, G. M. (2010). Netnography: A method specifically designed to study cultures and communities online. *Qualitative Report*, *15*(5), 1270–1275.
Bradshaw, S., & Howard, P. (2017). *Troops, Trolls and Troublemakers: A Global Inventory of Organized Social Media Manipulation* (Computational Propaganda Project Working Paper Series No. 2017.12) (p. 37). Oxford, England: University of Oxford. Retrieved from http://comprop.oii.ox.ac.uk/2017/07/17/troops-trolls-and-trouble-makers-a-global-inventory-of-organized-social-media-manipulation/
Bruner, J. (2016, June 15). Why 2016 is shaping up to be the Year of the Bot. Retrieved January 30, 2018, from https://www.oreilly.com/ideas/why-2016-is-shaping-up-to-be-the-year-of-the-bot
Burger. (2017, November 1). Interview Chicago Two.
Burrell, J. (2009). The Field Site as a Network: A Strategy for Locating Ethnographic Research. *Field Methods*, *21*(2), 181–199. https://doi.org/10.1177/1525822X08329699
Cancryn, A. (2016, November 9). Pro-Jeb strategist: "Data died tonight." Retrieved July 25, 2017, from http://politi.co/2flGzi7
Cassidy. (2016, September 16). Interview Seattle 4.
CBS 60 Minutes. (2016, November 12). Trump says his social media power key to victory. Retrieved May 22, 2017, from http://www.cbsnews.com/news/president-elect-trump-says-social-media-played-a-key-role-in-his-victory/
Chadwick, A. (2013). *The Hybrid Media System: Politics and Power*. Oxford, UK: Oxford University Press.
Chesney, R., Frauenfelder, M., Gentry, E., & Woolley, S. (2018, February). *Chat Apps: Frontiers and Challenges for Journalism*. Google News Lab and the Institute for the Future.
Cifuentes, J. (2016, July 19). Slack reveals the next 11 bot startups from its $80 million fund. Retrieved February 3, 2018, from https://venturebeat.com/2016/07/19/slack-reveals-the-first-11-bot-startups-from-its-80-million-fund/
Cohen, B. (1963). *Press and Foreign Policy*. Princeton, NJ: Princeton University Press.
Coldewey, D. (2012). Romney twitter account gets upsurge in fake followers, but from where. *NBC News*.
Coleman, G. (2014). *Hacker, Hoaxer, Whistleblower, Spy: The Many Faces of Anonymous*. London ; New York: Verso.
Collins, B. (2017, August 8). I Found Nicole Mincey, Trump's Biggest Twitter Fan. She Isn't a Bot, But She Has a Ton of Secrets. *The Daily Beast*. Retrieved from

https://www.thedailybeast.com/i-found-trump-fan-nicole-mincey-and-shes-not-a-bot
Coyne, B. (2017, November 7). How #Election2016 was Tweeted so far. Retrieved July 25, 2017, from https://blog.twitter.com/official/en_us/a/2016/how-election2016-was-tweeted-so-far.html
Craig, R. T. (1999). Communication Theory as a Field. *Communication Theory*, *9*(2), 119–161. https://doi.org/10.1111/j.1468-2885.1999.tb00355.x
Davis, D. K., & Jasinski, J. (1993). Beyond the Culture Wars: An Agenda for Research on Communication and Culture. *Journal of Communication*, *43*(3), 141–149. https://doi.org/10.1111/j.1460-2466.1993.tb01286.x
Delany, C. (2016, September 29). Why Is Trump Outspending Clinton Online? *Campaigns and Elections*. Retrieved from https://www.campaignsandelections.com/campaign-insider/why-is-trump-outspending-clinton-online
Diana. (2016, March 30). Interview New York 1.
DiResta, R., Little, J., Morgan, J., Neudert, L. M., Nimmo, B., & Woolley, S. (2017, November 2). The Bots That Are Changing Politics. Retrieved November 7, 2017, from https://motherboard.vice.com/en_us/article/mb37k4/twitter-facebook-google-bots-misinformation-changing-politics
Domi. (2016, October 27). Interview Seattle 2.
Dunn, J. (2016, November 9). Donal Trump Won Twitter, and That Was a Givaway that He Might Win the Presidency. *Business Insider*. Retrieved from http://www.businessinsider.com/donald-trump-twitter-popularity-chart-2016-11
Eastin, M. S., & LaRose, R. (2000). Internet Self-Efficacy and the Psychology of the Digital Divide. *Journal of Computer-Mediated Communication*, *6*(1), 0–0. https://doi.org/10.1111/j.1083-6101.2000.tb00110.x
Eastwood. (2016, May 14). Interview Seattle 6.
Ellyatt, H. (2017, November 8). How I helped get Trump elected: The president's digital guru Brad Parscale. Retrieved April 30, 2018, from https://www.cnbc.com/2017/11/08/how-i-helped-get-trump-elected-the-presidents-digital-guru-brad-parscale.html
Elmira. (2016, November 16). interview SF 1.
Farhi, P. (2009). The Twitter explosion: whether they are reporting about it, finding sources on it or urging viewers, listeners and readers to follow them on it, journalists just can't seem to get enough of the social networking service. Just how effective is it as a journalism tool? *American Journalism Review*, *31*(3), 26–32.
Ferrara, E., Varol, O., Davis, C., Menczer, F., & Flammini, A. (2016). The Rise of Social Bots. *Commun. ACM*, *59*(7), 96–104. https://doi.org/10.1145/2818717
Fisher, M., Cox, J. W., & Hermann, P. (2016, December 6). Pizzagate: From rumor, to hashtag, to gunfire in D.C. *Washington Post*. Retrieved from https://www.washingtonpost.com/local/pizzagate-from-rumor-to-hashtag-to-gunfire-in-dc/2016/12/06/4c7def50-bbd4-11e6-94ac-3d324840106c_story.html
Fisher, M., Woodrow Cox, J., & Hermann, P. (2017, December 6). Pizzagate: From rumor, to hashtag, to gunfire in D.C. Retrieved March 27, 2017, from

https://www.washingtonpost.com/local/pizzagate-from-rumor-to-hashtag-to-gunfire-in-dc/2016/12/06/4c7def50-bbd4-11e6-94ac-3d324840106c_story.html
Fleishman, C. (2016, October 24). Anti-Semitic Trump supporters found a chilling new way to attack Jewish journalists. *Mic*. Retrieved from https://mic.com/articles/157543/the-list-alt-right-donald-trump-trolls-harass-jewish-journalists-8chan-raid
Ford, H., Dubois, E., & Puschmann, C. (2016). Keeping Ottawa honest--one tweet at a time? Politicians, journalists, wikipedians, and their twitter bots. *International Journal of Communication*, 4891.
Forelle, M. C., Howard, P. N., Monroy-Hernandez, A., & Savage, S. (2015). *Political Bots and the Manipulation of Public Opinion in Venezuela* (SSRN Scholarly Paper No. ID 2635800). Rochester, NY: Social Science Research Network. Retrieved from https://papers.ssrn.com/abstract=2635800
Fox News. (2017, March 15). Exclusive: President Donald Trump. *Tucker Carlson Tonight*. New York, NY: Fox.
Franklin Fowler, E., Ridout, T. N., & Franz, M. M. (2017). Political Advertising in 2016: The Presidential Election as Outlier? *The Forum*, *14*(4), 445–469. https://doi.org/10.1515/for-2016-0040
Funke, D. (2017, December 14). Should we stop saying "fake news"? *Poynter*. Retrieved from https://www.poynter.org/news/should-we-stop-saying-fake-news
Geiger, R. S., & Ribes, D. (2011). Trace Ethnography: Following Coordination through Documentary Practices (pp. 1–10). Presented at the 44th Hawaii International Conference on System Sciences (HICSS), IEEE.
George. (2016, November 8). Interview New York 4.
Gillespie, T. (2011, October 19). Can an algorithm be wrong? Twitter Trends, the specter of censorship, and our faith in the algorithms around us. Retrieved September 13, 2015, from http://culturedigitally.org/2011/10/can-an-algorithm-be-wrong/
Gillespie, T. (2014). The Relevance of Algorithms. In T. Gillespie, P. J. Boczkowski, & K. A. Foot (Eds.), *Media Technologies: Essays on Communication, Materiality, and Society*. Cambridge, MA: MIT Press.
Goichman, R. (2017, February 15). Written by a Robot: Will Algorithms Kill Journalism? *Haaretz*. Retrieved from https://www.haaretz.com/israel-news/business/1.771758
Gonzales, H. M. S., & González, M. S. (2017). Bots as a news service and its emotional connection with audiences. The case of Politibot. *Doxa Comunicación. Revista Interdisciplinar de Estudios de Comunicación Y Ciencias Sociales*, *0*(25), 63–84.
Gorwa, R. (6/17). *Computational Propaganda in Poland: False Amplifiers and the Digital Public Sphere* (Computational Propaganda Working Paper Series No. 2017.4) (p. 38). Oxford, UK: Oxford Internet Institute, University of Oxford. Retrieved from http://comprop.oii.ox.ac.uk/2017/06/19/computational-propaganda-in-poland-false-amplifiers-and-the-digital-public-sphere/
Guilbeault, D., & Woolley, S. (2016, November 1). How Twitter Bots Are Shaping the Election. *The Atlantic*. Retrieved from https://www.theatlantic.com/technology/archive/2016/11/election-bots/506072/

Gupta, A., Lamba, H., Kumaraguru, P., & Joshi, A. (2013). Faking Sandy: Characterizing and Identifying Fake Images on Twitter During Hurricane Sandy. *Proceedings of the 22nd International Conference on the World Wide Web 2013*, 729–736.

Guy. (2015, August 20). Interview Online 4.

Harris, M. (2016, November 14). A Media Post-Mortem on the 2016 Presidential Election. Retrieved March 27, 2017, from http://www.mediaquant.net/2016/11/a-media-post-mortem-on-the-2016-presidential-election/

Harry. (2015, April 2). Seattle interview 1.

Harry. (2016, February 6). Interview Seattle 3.

Hazard Owen, L. (2016, February 5). The New York Times' new Slack 2016 election bot sends readers' questions straight to the newsroom. Retrieved February 3, 2018, from http://www.niemanlab.org/2016/02/the-new-york-times-launches-a-slack-2016-election-bot-that-accepts-questions-from-readers/

Herman, E. S., & Chomsky, N. (2002). *Manufacturing Consent: The Political Economy of the Mass Media* (Reprint edition). New York: Pantheon.

Higgins, A. (2016, May 30). Effort to Expose Russia's "Troll Army" Draws Vicious Retaliation. *New York Times*. Retrieved from https://www.nytimes.com/2016/05/31/world/europe/russia-finland-nato-trolls.html?_r=1

Hindman, M. (2008). *The Myth of Digital Democracy*. Princeton, NJ: Princeton University Press.

Hine, C. (2005). *Virtual Methods*. Oxford ; New York: Bloomsbury Academic.

Horst, H. A., & Miller, D. (2013). *Digital anthropology*. A&C Black. Retrieved from http://books.google.com/books?hl=en&lr=&id=9KuPzBgus7oC&oi=fnd&pg=PR5&dq=digital+anthropology&ots=uzgYGqVIHd&sig=OHQMy5JPTRN_zeAov_Jj0mi_3Gk

Howard, P. (2002). Network ethnography and the hypermedia organization: new media, new organizations, new methods. *New Media & Society*, *4*(4), 550–574.

Howard, P., Kollanyi, B., & Woolley, S. (2016). Bots and Automation over Twitter during the U.S. Election. *Oxford, UK: Project on Computational Propaganda*, (Data Memo 2016.4).

Howard, P. N. (2002). Network Ethnography and the Hypermedia Organization: New Media, New Organizations, New Methods. *New Media & Society*, *4*(4), 550–574. https://doi.org/10.1177/146144402321466813

Howard, P. N. (2011). *Castells and the Media: Theory and Media*. Polity.

Howard, P. N., & Hussain, M. M. (2013). *Democracy's Fourth Wave?: Digital Media and the Arab Spring*. Oxford University Press. Retrieved from http://books.google.com/books?hl=en&lr=&id=ayHOyrmmT8kC&oi=fnd&pg=PP2&dq=howard+digital+activism+2013&ots=mJDwHv1nfa&sig=jvnGD3UO3Bg1GuzR2wayT4WD8pI

Howard, P. N., Kollanyi, B., Bolsover, G., Bradshaw, S., & Neudert, L.-M. (2017). Junk News and Bots during the U.S. Election: What Were Michigan Voters Sharing Over Twitter? Retrieved from comprop.oii.ox.ac.uk

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

Howard, P. N., Woolley, S., & Calo, R. (2018). Algorithms, bots, and political communication in the US 2016 election: The challenge of automated political communication for election law and administration. *Journal of Information Technology & Politics*, *0*(0), 1–13. https://doi.org/10.1080/19331681.2018.1448735

Hwang, T., Pearce, I., & Nanis, M. (2012). Socialbots: Voices from the fronts. *Interactions*, *19*(2), 38–45.

Jam. (2016, August 18). Interview Moorend 2.

Jessie. (2016, March 7). Interview Detroit 2.

Johnson, L. (2016, September 21). More Than a Third of Trump's and Clinton's Twitter Followers Are Reportedly Fake. Retrieved July 25, 2017, from http://www.adweek.com/digital/more-third-presidential-candidates-twitter-followers-are-reportedly-fake-173628/

Jose. (2016, December 1). Interview New York 6.

Karpf, D. (2012). *The MoveOn effect: The unexpected transformation of American political advocacy*. Oxford University Press. Retrieved from http://books.google.com/books?hl=en&lr=&id=YReQ5quOFvUC&oi=fnd&pg=PP2&dq=The+moveon+effect&ots=GK-Vxi9cUm&sig=YfKIp3HYzKlIorhd4MjeEMjKbNU

Keefe, J. (2017, October 3). Announcing Quackbot, a Slack bot for journalists from Quartz and DocumentCloud. Retrieved November 7, 2017, from https://bots.qz.com/2017/10/03/announcing-quackbot-a-slack-bot-for-journalists-from-quartz-and-documentcloud/

Keohane, J. (2017, February 16). A Robot May Have Written This Story. Retrieved November 7, 2017, from https://www.wired.com/2017/02/robots-wrote-this-story/

Kollanyi, B., Howard, P. N., & Woolley, S. C. (2016). *Bots and Automation over Twitter during the U.S. Election* (Data Memo No. 2016.4) (p. 5). Oxford, UK: Project on Computational Propaganda. Retrieved from http://www.politicalbots.org

König, J. (2016, December 18). We built a Twitter bot that replies to people who share fake news. Retrieved November 7, 2017, from https://medium.com/thoughts-on-journalism/we-built-a-twitter-bot-that-replies-to-people-who-share-fake-news-d23127c1ed15

Kreiss, D. (2016). *Prototype Politics: Technology-Intensive Campaigning and the Data of Democracy*. New York, NY: Oxford University Press.

Kreiss, D. (2017, May). *Digital Campaigns, Social Media and Elections: New Platforms, New Research*. Presented at the International Communication Association, San Diego, CA.

Lapowsky, I. (2016, November 15). This Is How Facebook Actually Won Trump the Presidency. Retrieved March 27, 2017, from https://www.wired.com/2016/11/facebook-won-trump-election-not-just-fake-news/

Larsson, A. O., & Hallvard, M. (2015). Bots or journalists? News sharing on Twitter. *Communications*, *40*(3), 361–370. https://doi.org/10.1515/commun-2015-0014

Latour, B., & Venn, C. (2002). Morality and Technology The End of the Means. *Theory, Culture & Society*, *19*(5-6), 247–260. https://doi.org/10.1177/026327602761899246
Lauter, D. (2016, November 8). One last look at the polls: Hillary Clinton's lead is holding steady. Retrieved April 29, 2018, from /nation/politics/trailguide/la-na-election-day-2016-a-last-look-at-the-polls-clinton-lead-1478618744-htmlstory.html
Leah. (2017, March 29). Interview New York 7.
Leo. (2016, March 8). Interview Detroit 1.
Lichterman, J. (2016, October 20). This is how The New York Times is using bots to create more one-to-one experiences with readers. Retrieved November 7, 2017, from http://www.niemanlab.org/2017/04/this-is-how-the-new-york-times-is-using-bots-to-create-more-one-to-one-experiences-with-readers/
Lloyd, A. (2016, February 25). Our friends, the bots? *Data & Society: Points*. Retrieved from https://points.datasociety.net/our-friends-the-bots-34eb3276ab6d
Lokot, T., & Diakopoulos, N. (2016). News Bots: Automating News and Information Dissemination on Twitter. *Digital Journalism*, *4*(6), 682–699. https://doi.org/10.1080/21670811.2015.1081822
Maggie. (2014, May 13). Interview Budapest 1.
Manjoo, F. (2017, May 31). How Twitter Is Being Gamed to Feed Misinformation. *The New York Times*. Retrieved from https://www.nytimes.com/2017/05/31/technology/how-twitter-is-being-gamed-to-feed-misinformation.html
Markham, A. N., & Baym, N. K. (2009). *Internet inquiry: conversations about method*. Los Angeles: Sage Publications.
Mayer. (2016, August 14). Interview Moorend 1.
McCombs, M. E., & Shaw, D. L. (1972). The Agenda Setting Function of Mass Media. *Public Opinion Quarterly*, *36*(2), 176–187. https://doi.org/10.1086/267990
McCombs, M. E., Shaw, D. L., & Weaver, D. H. (1997). *Communication and Democracy: Exploring the Intellectual Frontiers in Agenda-setting Theory*. London, UK: Psychology Press.
Mel. (2017, February 5). Interview Brooklyn 1.
Metaxas, P. T., & Mustafaraj, E. (2012a). Social media and the elections. *Science*, *338*(6106), 472–473.
Metaxas, P. T., & Mustafaraj, E. (2012b). Social media and the elections. *Science*, *338*(6106), 472–473.
Monaco, N., & Nyss, C. (2018, February). *State Sponsored Trolling: How Governments are Deploying Fake News as Part of Broader Harrassment Campaigns*. Institute for the Future Working Research Papers.
Mullin, B. (2016, July 17). BuzzFeed's newest political reporter is a bot. Retrieved November 7, 2017, from https://www.poynter.org/news/buzzfeeds-newest-political-reporter-bot
Musgrave, S. (2017, August 9). "I Get Called a Russian Bot 50 Times a Day." *POLITICO Magazine*. Retrieved from http://politi.co/2viAxZA

Mustafaraj, E., & Metaxas, P. T. (2010). From obscurity to prominence in minutes: Political speech and real-time search. Presented at the WebSci10. Retrieved from http://repository.wellesley.edu/computersciencefaculty/9/

Navi. (2017, February 4). Interview New York 5.

Neff, Gi. (2012). *Venture Labor: Work and the Burden of Risk in Innovative Industries*. Cambridge, Mass: The MIT Press.

Neff, G., & Nagy, P. (2016a). Automation, Algorithms, and Politics| Talking to Bots: Symbiotic Agency and the Case of Tay. *International Journal of Communication*, *10*(0), 17.

Neff, G., & Nagy, P. (2016b). Talking to bots: Symbiotic agency and the case of tay. *International Journal of Communication*, *10*(0), 17.

Olaf. (2016, September 18). Interview Seattle 5.

Ott, B. L. (2017). The age of Twitter: Donald J. Trump and the politics of debasement. *Critical Studies in Media Communication*, *34*(1), 59–68. https://doi.org/10.1080/15295036.2016.1266686

Owen, T. (2015). *Disruptive Power: The Crisis of the State in the Digital Age*. Oxford, UK: Oxford University Press.

Oxford Dictionaries. (2017). constituent. *Oxford Dictionaries*. Oxford, UK. Retrieved from https://en.oxforddictionaries.com/definition/constituent

Pan. (2017, March 8). Interview Chicago 1.

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

Patterson, T. E. (2016). News Coverage of the 2016 Election: How the Press Failed Voters. *Shorestein Center on Media, Politics, and Public Policy*.

Pazzanese, C. (2017, March 30). The Puzzle in Politics and Polling. *Harvard Gazette*. Retrieved from http://news.harvard.edu/gazette/story/2017/03/nate-silver-says-conventional-wisdom-not-data-killed-2016-election-forecasts/

Pelly, S. (2017, March 26). How fake news becomes a popular, trending topic. *60 Minutes*. New York, NY: CBS. Retrieved from https://www.cbsnews.com/news/how-fake-news-find-your-social-media-feeds/

Phelps, A. (2017, April 26). This is how The New York Times is using bots to create more one-to-one experiences with readers. Retrieved November 7, 2017, from http://www.niemanlab.org/2017/04/this-is-how-the-new-york-times-is-using-bots-to-create-more-one-to-one-experiences-with-readers/

PJNet. (n.d.). Patriot Journalist Network. Retrieved October 26, 2017, from http://patriotjournalist.com/

Proxy. (2016). *Merriam-Webster*. Retrieved from https://www.merriam-webster.com/dictionary/proxy

Ratkiewicz, J., Conover, M., Meiss, M., Goncalves, B., Flammini, A., & Menczer, F. (2011). Detecting and Tracking Political Abuse in Social Media. In *ICWSM*. Retrieved from http://www.aaai.org/ocs/index.php/ICWSM/ICWSM11/paper/viewFile/2850/3274

Ratkiewicz, J., Conover, M., Meiss, M., Goncalves, B., Patil, S., Flammini, A., & Menczer, F. (2011). Truthy: mapping the spread of astroturf in microblog streams. In *Proceedings of the 20th international conference companion on World wide web* (pp. 249–252). ACM. Retrieved from http://dl.acm.org/citation.cfm?id=1963301

Read, M. (2016, November 9). Donald Trump Won Because of Facebook. *New York Magazine*. Retrieved from http://nymag.com/selectall/2016/11/donald-trump-won-because-of-facebook.html

Reke. (2016, September 27). Interview New York 2.

Robb, A. (2017, November 16). Pizzagate: Anatomy of a Fake News Scandal. *Rolling Stone*. Retrieved from https://www.rollingstone.com/politics/news/pizzagate-anatomy-of-a-fake-news-scandal-w511904

Robb, J. (2007, August 23). When Bots Attack. *WIRED*, *15*(09). Retrieved from https://www.wired.com/2007/08/ff-estonia-bots/

Rockefeller, H. (2011, February 16). UPDATED: The HB Gary Email That Should Concern Us All. Retrieved February 24, 2014, from http://www.dailykos.com/story/2011/02/16/945768/-UPDATED-The-HB-Gary-Email-That-Should-Concern-Us-All

Roger. (2016, April 18). Interview New York 3.

Schauer, F. (1978). Fear, Risk and the First Amendment: Unraveling the Chilling Effect. *Boston University Law Review*, *58*, 685.

Schramm, W. (1983). The Unique Perspective of Communication: A Retrospective View. *Journal of Communication*, *33*(3), 6–17. https://doi.org/10.1111/j.1460-2466.1983.tb02401.x

Schreckinger, B. (2016a, September 30). Inside Trump's "cyborg" Twitter army. Retrieved May 23, 2017, from http://politi.co/2dyhCD0

Schreckinger, B. (2016b, September 30). Inside Trump's "cyborg" Twitter army. *POLITICO*. Retrieved from http://politi.co/2dyhCD0

Segerberg, A., & Bennett, W. L. (2011). Social Media and the Organization of Collective Action: Using Twitter to Explore the Ecologies of Two Climate Change Protests. *The Communication Review*, *14*(3), 197–215. https://doi.org/10.1080/10714421.2011.597250

Seligman, A. B. (1995). *The Idea of Civil Society*. Princeton University Press.

Seward, Z. (2016, November 29). Why we're starting the Quartz Bot Studio. Retrieved November 7, 2017, from https://blog.qz.com/why-were-starting-the-quartz-bot-studio-ae535cc9ed45

Shao, C., Ciampaglia, G., Varol, O., Flammini, A., & Menczer, F. (2017). The Spread of Misinformation by Social Bots. *ArXiv*.

Shelon. (2016, February 2). Interview Seattle 1.

Siri. (2016, May 1). East Coast Interview 1.

Stelter, B. (2017, January 16). Donald Trump's Twitter milestone: 20 million followers. Retrieved July 25, 2017, from http://money.cnn.com/2017/01/16/media/donald-trump-twitter-20-million-followers/index.html

Sykes, C. J. (2017, February 4). Why Nobody Cares the President Is Lying. *The New York Times*. Retrieved from

https://www.nytimes.com/2017/02/04/opinion/sunday/why-nobody-cares-the-president-is-lying.html
Tanaka, K. (2017, July 20). The Threat of Computational Propaganda. *Nightly News*. Tokyo, JP: NHK. Retrieved from http://www3.nhk.or.jp/news/web_tokushu/2017_0720.html?utm_int=news_content
Tomas. (2017, September 29). Interview Online 3.
Tufekci, Z., & Wilson, C. (2012). Social Media and the Decision to Participate in Political Protest: Observations From Tahrir Square. *Journal of Communication*, *62*(2), 363–379. https://doi.org/10.1111/j.1460-2466.2012.01629.x
Verkamp, J. P., & Gupta, M. (2013). Five Incidents, One Theme: Twitter Spam as a Weapon to Drown Voices of Protest. Presented at the USENIX Workshop on Free and Open Communications on the Internet (FOCI), Washington, D.C.
Vertesi, J. (2008). Mind the Gap The London Underground Map and Users' Representations of Urban Space. *Social Studies of Science*, *38*(1), 7–33. https://doi.org/10.1177/0306312707084153
Vertesi, J. (2015). *Seeing Like a Rover: How Robots, Teams, and Images Craft Knowledge of Mars*. Chicago ; London: University Of Chicago Press.
Vincent, J. (2016, March 24). Twitter taught Microsoft's friendly AI chatbot to be a racist asshole in less than a day. *The Verge*. Retrieved from https://www.theverge.com/2016/3/24/11297050/tay-microsoft-chatbot-racist
Vis, F. (2013). Twitter as a Reporting Tool for Breaking News. *Digital Journalism*, *1*(1), 27–47. https://doi.org/10.1080/21670811.2012.741316
Walker, A. (2014, March 19). Quakebot: An Algorithm That Writes the News About Earthquakes. *Gizmodo*. Retrieved from https://gizmodo.com/quakebot-an-algorithm-that-writes-the-news-about-earth-1547182732
Walter. (2016, October 24). Interview Online 1.
Wang, S. (2016, October 29). Quartz launches its Bot Studio with $240K from Knight, and plans for Slack and Echo » Nieman Journalism Lab. *BuzzFeed*. Retrieved from http://www.niemanlab.org/2016/11/quartz-launches-its-bot-studio-with-a-quarter-million-from-knight-and-plans-for-slack-and-amazon-echo/
Wang, Y., Li, Y., & Luo, J. (2016). Decifering the 2016 U.S. Presidential Campaign in the Twitter Sphere: A Comparison of Trumpists and Clintonists. In *Proceedings of the Tenth International AAAI Conference on Web and Social Media*. Cologne, DE: AAAI. Retrieved from https://www.aaai.org/Press/Proceedings/icwsm16.php
Wardle, C. (2017, December). Disinformation gets worse. Retrieved May 3, 2018, from http://www.niemanlab.org/2017/12/disinformation-gets-worse/
Washington Post Staff. (2017, May 8). Full transcript: Sally Yates and James Clapper testify on Russian election interference. *Washington Post*. Retrieved from https://www.washingtonpost.com/news/post-politics/wp/2017/05/08/full-transcript-sally-yates-and-james-clapper-testify-on-russian-election-interference/
WashPostPR. (2016, October 20). The Washington Post's bot for Messenger to capture how people feel the final days of the 2016 election. *Washington Post*. Retrieved from https://www.washingtonpost.com/pr/wp/2016/10/20/the-washington-posts-bot-for-messenger-to-capture-how-people-feel-the-final-days-of-the-2016-election/

Weaver, J. F. (2015, February 25). Who's responsible when a twitter bot sends a threatening tweet? *Slate*. Retrieved from http://www.slate.com/blogs/future_tense/2015/02/25/who_is_responsible_for_death_threats_from_a_twitter_bot.html

Wehner, P. (2016, January 14). Why I Will Never Vote for Donald Trump. *The New York Times*. Retrieved from https://www.nytimes.com/2016/01/14/opinion/campaign-stops/why-i-will-never-vote-for-donald-trump.html

Wesleyan Media Project. (n.d.). Political Advertising in 2016: The Presidential Election as an Outlier? Retrieved from http://mediaproject.wesleyan.edu/category/releases/

Williams, B. A., & Carpini, M. X. D. (2011). *After Broadcast News: Media Regimes, Democracy, and the New Information Environment*. New York, NY: Cambridge University Press.

Woolley, S. (2016). Automating power: Social bot interference in global politics. *First Monday*, *21*(4). https://doi.org/10.5210/fm.v21i4.6161

Woolley, S., boyd, danah, Broussard, M., Elish, M., Fader, L., Hwang, T., … Shorey, S. (2016, February 23). How to Think About Bots. Retrieved January 9, 2017, from http://motherboard.vice.com/read/how-to-think-about-bots

Woolley, S. C. (2015, August 4). #HackingTeam Leaks: Ecuador is Spending Millions on Malware, Pro-Government Trolls. Retrieved July 24, 2017, from https://advox.globalvoices.org/2015/08/04/hackingteam-leaks-ecuador-is-spending-millions-on-malware-pro-government-trolls/

Woolley, S. C. (2016). Automating power: Social bot interference in global politics. *First Monday*, *21*(4). Retrieved from http://firstmonday.org/ojs/index.php/fm/article/view/6161

Woolley, S. C., & Howard, P. N. (6/17). *Computational Propaganda Worldwide: Executive Summary* (Computational Propaganda Working Paper Series No. 2017.11) (p. 15). Oxford, UK: Oxford Internet Institute, University of Oxford. Retrieved from http://comprop.oii.ox.ac.uk/2017/07/11/computational-propaganda-worldwide-executive-summary/

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

Woolley, S. C., & Howard, P. N. (2016b). Introduction: Political communication, computational propaganda, and autonomous agents. *International Journal of Communication*, *10*(0), 9.

Woolley, S., & Guilbeault, D. (2017, April 10). Bots aren't just service tools—they're a whole new form of media. Retrieved May 13, 2017, from https://qz.com/954255/bots-are-the-newest-form-of-new-media/

Ziggy. (2016, August 21). Interview Moorend 3.