Introduction: What Is Computational Propaganda?

Digital technologies hold great promise for democracy. Social media tools and the wider resources of the Internet offer tremendous access to data, knowledge, social networks, and collective engagement opportunities, and can help us to build better democracies (Howard, 2015; Margetts et al., 2015). Unwelcome obstacles are, however, disrupting the creative democratic applications of information technologies (Woolley, 2016; Gallacher et al., 2017; Vosoughi, Roy, & Aral, 2018). Massive social platforms like Facebook and Twitter are struggling to come to grips with the ways their creations can be used for political control. Social media algorithms may be creating echo chambers in which public conversations get polluted and polarized. Surveillance capabilities are outstripping civil protections. Political "bots" (software agents used to generate simple messages and "conversations" on social media) are masquerading as genuine grassroots movements to manipulate public opinion. Online hate speech is gaining currency. Malicious actors and digital marketers run junk news factories that disseminate misinformation to harm opponents or earn click-through advertising revenue.

It is no exaggeration to say that coordinated efforts are even now working to seed chaos in many political systems worldwide. Some militaries and intelligence agencies are making use of social media as conduits to undermine democratic processes and bring down democratic institutions altogether (Bradshaw & Howard, 2017). Most democratic governments are preparing their legal and regulatory responses. But unintended consequences from over-regulation, or regulation uninformed by systematic research, may be as damaging to democratic systems as the threats themselves.

We live in a time of extraordinary political upheaval and change, with political movements and parties rising and declining rapidly (Kreiss, 2016; Anstead, 2017). In this fluctuating political environment, digital technologies provide the platform for a great deal of contemporary civic engagement and political action (Vaccari, 2017). Indeed, a large amount of research has shown that social media play an important role in the circulation of ideas and conversation about politics and public policy. Increasingly, however, social media platforms are also vehicles for manipulative disinformation campaigns. Political campaigns, governments, and regular citizens around the world are employing combinations of people and bots—automated software built to mimic real "real users—in an attempt to artificially shape public life (Woolley, 2016; Gallacher et al., 2017). But there are still open, and difficult to answer, questions about the specific mechanisms of influence for particular voters, and how governments, news organizations, and civil society groups should respond. How do new forms of civic engagement affect political outcomes? To what extent do online echo chambers and selective exposure to information promote political extremism? How can civil activists respond effectively to "trolling" by hostile political agents?

Computational propaganda is a term that neatly encapsulates this recent phenomenon—and emerging field of study—of digital misinformation and manipulation. As a communicative practice, computational propaganda describes the use of algorithms, automation, and human curation to purposefully manage and distribute misleading information over social media networks (Woolley & Howard, 2016a). As part of the process, coders and their automated software products (including bots) will learn from and imitate legitimate social media users in order to manipulate public opinion across a diverse range of platforms and device networks. These bots are built to behave like real people (for example, automatically generating and responding to conversations online) and then let loose over social media sites in order to amplify or suppress particular political messages. These "automated social actors" can be used to bolster particular politicians and policy positions—supporting them actively and enthusiastically, while simultaneously drowning out any dissenting voices (Abokhodair, Yoo, & McDonald, 2015). They can be managed in conjunction with human troll armies to

"manufacture consensus" or to otherwise give the illusion of general support for a (perhaps controversial) political idea or policy, with the goal of creating a bandwagon effect (Woolley & Guilbeault, 2017). Computational propaganda therefore forms part of a suite of dubious political practices that includes digital astroturfing, state-sponsored trolling, and new forms of online warfare known as PsyOps or InfoOps wherein the end goal is to manipulate information online in order to change people's opinions and, ultimately, behavior.

However, trying to understand computational propaganda only from a technical perspective—as a set of variables, models, codes, and algorithms—plays into the hands of those who create it, the platforms that serve it, and the firms that profit from it (Bolsover & Howard, 2017). The very act of describing something as purely "technical" or in very mechanistic terms may make it seem unbiased and inevitable. This is clearly a dangerous position to take, and we must look to the emerging discipline of "social data science" to help us understand the complex socio-technical issues at play, and the influence of technology (including computational propaganda) on politics. As part of this process, social data science researchers must maintain a critical stance toward the data they use and analyze, so as to ensure that they are critiquing as they go about describing, predicting, or recommending changes in the way technology interacts with our political systems. If academic research on computational propaganda does not engage fully with the systems of power and knowledge that produce it (that is, the human actors and motivations behind it), then the very possibility of improving the role of social media platforms in public life evaporates (Bolsover, 2017). We can only hope to understand and respond appropriately to a problem like computational propaganda's impact on our political systems by undertaking computational research alongside qualitative investigation—by addressing the computational as well as the political.

Computational propaganda, with this in mind, can therefore be understood to incorporate two important components: the technical and the social. As a technical phenomenon, we can define computational propaganda as the assemblage of social media platforms, autonomous agents, algorithms, and big data tasked with the manipulation of public opinion (Woolley & Howard, 2016b). "Computational" propaganda is of course a recent form of the propaganda that has existed in our political systems for millennia—communications that deliberately subvert symbols, appealing to our baser emotions and prejudices and bypassing rational thought, to achieve the specific goals of its promoters—with computational propaganda understood as propaganda created or disseminated by computational means. Automation, scalability, and anonymity are hallmarks of computational propaganda. The pernicious advantage of computational propaganda is in enabling the rapid distribution of large amounts of content, sometimes personalized in order to fool users into thinking that messages originate in their extended network of family and friends. In this way, computational propaganda typically involves one or more of the following ingredients: bots that automate content delivery; fake social media accounts that require some (limited) human curation; and junk news—that is, misinformation about politics and public life.

The political bots we have already mentioned as being integral to the spread of computational propaganda are software programs or agents that are created to perform simple, repetitive, typically text-based tasks. Generally speaking, bots are used to computationally enhance the ability of humans to get work done online, both in terms of volume and speed. This work can be benign and extremely useful: most of the internal links that allow us to navigate Wikipedia are created and maintained by bots. When bots are programmed with human attributes or abilities—in order to pass as genuine social media users, for instance—they are referred to as social bots or chat bots. They can be used to perform mundane tasks like gathering information, but they can also interact with people and systems. This could involve simple tasks like delivering news and information—automated updates about the weather, sports news, and share values, for example. They can also be used to support more malicious activities, such as spamming and harassment. But regardless of whether they are put to a

benign or malicious task, they are able to rapidly deploy messages, interact with other users' content, and even affect or manipulate trending algorithms—all while passing as human users. Political bots—that is, social bots used for political manipulation—thus represent an effective tool for driving online propaganda and hate campaigns. One person, or a small group of people, can fairly easily create and coordinate an army of political bots on Twitter, YouTube, or Instagram to give the illusion of a large-scale consensus or interest in a particular issue.

Governments and political actors around the world have used political bots—programmed to appear and behave like genuine citizens—to drown out and harass the opposition and to push their own messages. Political campaigns (and their civilian supporters) have deployed political bots and computational propaganda during recent elections in order to swing the vote, or to defame and intimidate the opposition. Anonymous political actors have spread false news reports, and coordinated disinformation campaigns and troll mobs to attack human rights defenders, civil society groups, and journalists. Computational propaganda is an extremely powerful new communication tool—and it is being used against democratic actors and institutions worldwide.

Automation and Algorithms as Tools for Political Communication

The term computational propaganda can be used to describe the recent series of digital attacks on civic society. The "computational" part of the equation is an important one. Data-driven techniques and tools like automation (bots) and algorithms (decision-making code) allow small groups of actors to megaphone highly specific, and sometime abusive and false, information into mainstream online environments. Rapid cycles of sharing, repurposing, and further dissemination often ensue.

During the 2016 US presidential election, for instance, far-right users on the 8chan imageboard spread a meme featuring Hillary Clinton, the Star of David, and a background of money. This image was then disseminated on sites like Facebook and subsequently shared and re-shared by mainstream conservatives. Presidential candidate Donald Trump then re-tweeted the image. The media picked up on the massive uptick in online chatter and began writing stories on the subject. The tactic—using hate and the viral spread of disinformation to undermine opposition—is not necessarily an new one. Russian propagandists and others have made healthy use of it in recent years (Castle, 2015). However, the rate at which this information is now routinely seeded—and the degree of confusion created by its rapid growth and spread online—is new.

The history of computational propaganda is of course brief, relative to the much longer history of traditional forms of political propaganda. But over the last six years state and nonstate political actors—from candidates for office to hacking collectives—have successfully swayed opinion and behavior during critical elections, security crises, and other important political events (Woolley, 2016). Powerful (and often anonymous) political actors have used computational propaganda techniques to perpetrate political attacks, to spread disinformation, censor and attack journalists, and create fake trends. In the last five years clear-cut cases of this have been observed in Argentina, Australia, Azerbaijan, Bahrain, Brazil, China, Iran, Italy, Mexico, Russia, South Korea, Saudi Arabia, Turkey, the United Kingdom, the United States, and Venezuela.

Automation and anonymity lie at the heart of computational propaganda, and underpin what is both interesting and important about it as a new field of academic enquiry. "Computational" doesn't just mean that these acts of persuasion happen on a computer or online. Rather, it underscores the fact that these political strategies rely on computational enhancement. Automation allows propaganda attacks to be scaled. Anonymity allows perpetrators to remain unknown. In 2015, the security firm Incapsula published a study finding that bots generate

almost half of all Web traffic—an extraordinary proportion (Zeifman, 2015). Within the social media sphere, estimates suggest that over a third of Twitter's users are in fact bots—automated software-driven accounts built to pass as real people. Estimates claim that within two years bots will generate around 10 percent of all activity on popular social media sites. These broad estimates of bot activity might sound rather horrifying (if we value social media as a space of genuine, unmediated connection with other people), but the details are even more so—many bots now maintain a parallel presence on several social media sites concurrently, to lend themselves an aura of human credibility. They also mimic human lifestyles—adhering to a believable sleep-wake cycle, making them harder to identify based on usage patterns alone.

Social bots on dating apps like Tinder are programmed to not respond immediately to human advances, but to delay their response as a human might (Melendez, 2015). Indeed, as has been pointed out by Chu et al. (2010), on Twitter a bot can do nearly everything a human can do through the Twitter API, and they note the ever-increasing difficulty of distinguishing between scripts generated by humans, bots or cyborgs (that is, a bot-aided human, or a human-aided bot).

Many bots are launched via a social media platform's application programming interface (API). Some sites, like Twitter, have more open APIs and less strict policies around bot use. On Twitter, bots can be directly plugged into one of many APIs. They can process information and activity on Twitter in real time, and respond to any comments and users that are relevant to their script (for example, that are identifiable as promoting or following a particular user or view). Facebook has more stringent policies about the use of automation, and maintains a "real name" policy that requires every user to verify their (unique, human) identity, but it still has problems with manipulative or otherwise problematic automation. In fact, in 2012 Facebook publicly announced that it intended to combat the fake accounts present on the social network, which amounted to 8.7 percent of all accounts (Wasserman, 2012). This percentage might seem small at first glance, but it represented 83 million accounts; equivalent to the entire population of Germany. Facebook was explicit in stating that these fake accounts and their "fraudulent likes" were antithetical to its purpose:

> Real identity, for both users and brands . . . is important to not only Facebook's mission of helping the world share, but also the need for people and customers to authentically connect to the Pages they care about . . . Facebook was built on the principle of real identity and we want this same authenticity to extend to Pages. We undoubtedly expect that this will be a positive change for anyone using Facebook. ("Improvements To Our Site Integrity Systems," 2012)

However, Facebook—in common with all the world's most popular social media platforms—continues to struggle with the bot problem.
In 2014 a bot named "Eugene Goostman" passed the Turing Test for the first time; meaning that it fooled a third of the judges into believing mistakenly that it was human, following a five-minute conversation between bot and judge (Aamoth, 2014). Bots are becoming increasingly humanlike in their speech and behavior, and thus more difficult to detect. They can be bought cheaply, with armies of bots built to like particular content or send message "bombs" costing less than 100 US dollars.

The Social Data Science of Political Communication

Algorithms and other computational tools now play an important political role in areas like news consumption, issue awareness, and cultural understanding (Gillespie, 2012; Sandvig et al., 2016)—leading to concern within the social sciences, especially within media studies and science and technology studies, about their impact on social life. The various problems thrown

up by this intersection are also explored in the information and computer sciences literature. Working in conversation with research from the computer sciences (Mitter, Wagner, & Strohmaier, 2014; Ferrara et al., 2016), communication and media oriented work has shown that political actors around the globe are using social media bots in efforts to both facilitate and control communication (Woolley & Howard, 2016a, 2016b; Woolley, 2016). Bots have been used by political campaigns and candidates in order to manipulate public opinion by disrupting activist attempts to organize, and also to create the illusion of popularity and consensus. This work highlights the increasing sophistication of modern social bots, and also their potential to threaten civic life both online and offline.

One particularly damaging form of computational propaganda is false news reports, widely distributed by bots over social media platforms like Twitter, Facebook, Reddit, and beyond. These social media platforms have served significant volumes of fake, sensational, and other forms of "junk news" during sensitive political moments over the last several years. However, most platforms reveal little about how much of this content there is, or what its impact on users may be. But in a marker of how important this problem might actually be, the World Economic Forum recently identified the rapid spread of misinformation online as among the top 10 perils to society (World Economic Forum, 2014). Previous research has found that social media favors sensationalist content, regardless of whether the content has been fact-checked or is from a reliable source (Vicario et al., 2016). When distribution of this junk news is backed by automation, either via political bots or through the platform operator's own dissemination algorithms, political actors have a powerful set of tools for computational propaganda. Both state and nonstate political actors can deliberately manipulate and amplify nonfactual information online, to their own ends.

Building and using bots has also been discussed as a kind of "agnostic commercialism." Bot builders act as hired guns, selling bots on freelancing platforms like Fiverr, with little concern for how they will be used or by whom. Using bots to spread advertisements or to attack opponents online has a long history, existing in email and other chat mediums before spreading to social media platforms like Twitter and Facebook. These newer forms can gather information on users in order to push a particular argument or agenda, often via hashtags (Hwang, Pearce, & Nanis, 2012). Experiments on the efficacy of such bots have shown that they can infiltrate social networks on sites like Facebook with a high degree of success, and that they can bypass the security systems intended to protect users from just such attacks (Boshmaf et al., 2013). In our own interviews with programmers who build and deploy such bots, many have told us that their work is purely mercenary—that they are apolitical in their views, and driven solely by a desire to make money online.

Of course, voter manipulation existed long before bots became mainstream on social media. Over a decade ago, Howard (2005) established the study of political "astroturf" movements, defining astroturfing as the process of seeking electoral victory or legislative relief for grievances by helping political actors find and mobilize a sympathetic public using the Internet. This campaign strategy can be used to create the image of public consensus where there is none, or to give a false impression of the popularity of a candidate or public policy idea. Almost as soon as social media took off as arguably the most important means of receiving news and communicating with our peers, network automation was used to support political communication in similar ways. Ratkiewicz et al. (2011) examined the ways in which Twitter bots were deployed before and during the 2010 US midterm congressional campaigns. They explored social bot-driven attacks upon candidates for the House and Senate, and suggested that these technological actors formed part of larger "astroturf" political efforts. Social bots, or "sock puppets," were harnessed in this context for their anonymity and ubiquity.

While there is a great deal of academic work exploring the tremendous democratic potential of the Internet, recent research has also shown how the liberating uses of the Internet can be

compromised when governing elites use them as tools for social control. Within- and between-country digital divides may also have an impact on how social media are used in public life—analysis by Schradie (2011) suggests that there is a class-based gap between producers of online content, and consumers. Still others argue that the now widespread political normalization of social media systems has allowed the politically powerful to leverage these tools for coercion and control (Karpf, 2012). Indeed, states that have exercised firm control over their own Internet development from the beginning—such as China, Singapore, and Iran—have proven success in online control (Kalathil & Boas, 2010).

A Mixed-Method Approach for Understanding Computational Propaganda

Media scholars have been concerned with the study of propaganda, and its effect upon public opinion, at least since the seminal work of Lazarsfeld (1941) and Lasswell (1948). Our own multi-year study of computational propaganda worldwide, which we present in this edited volume, picks up this line of research in order to understand the ways in which algorithms, automation, and social media are being used to promote the messages of political actors in many different kinds of democratic and authoritarian regimes around the world. The case studies we present all begin with a basic set of research questions crafted for comparability. Does computational propaganda occur as part of a country's current political landscape? What are its forms, types, or styles? What is its impact on public life? Each case study also considers the impact of the observed phenomena on the country's political institutions. How might political bot activity run afoul of its election law? Which computational propaganda campaigns had a significant impact, and how might they be prevented in the future?"

In keeping with our previous point about the importance of tackling this area from both the technical and social side, the findings we present are the result of knowledge generated via multiple social and data science methods. Our research team made use of both qualitative and quantitative methods of analysis. This mixed-method approach enables the case studies to speak to concerns at the intersection of several disciplines, especially those focused on social science, law, and computer science. We have conducted qualitative and quantitative content analysis of news coverage about computational algorithms. We have performed big data analysis of large networks of Facebook, Twitter, and Weibo users. Researchers have used a variety of methods in cataloguing their country-specific case studies including, but not limited to interviews with the victims of attacks, interviews with those who have worked to produce political bots and social media-based propaganda, process tracing, participant observation, and social network analysis. Each case required different approaches and tools.

The research team involved 12 researchers across nine countries who, altogether, interviewed 65 experts, analyzed tens of millions of posts on seven different social media platforms during scores of elections, political crises, and national security incidents. The[…]"

Computational Social Science. The research team applied a suite of machine learning techniques, including regression analysis, k-core analysis, and topic discovery methods in order to analyze public data collected from social networking sites, from surveys, and interviews. The goal of these methods is usually to map social influence, identify region-specific hot button issues that polarize social groups—like race, religion, immigration, and gender—and to track the use of online misinformation campaigns to influence voters.

Qualitative Ethnography, Participant Observation, and Fieldwork. Spending time with the designers of social media platforms or political communication experts yielded significant insights into how the affordances of these technical systems can limit citizens' choices for self-expression. Systematic interviews with political consultants, data mining firms, and civil society victims of attacks, reveal much about the economic incentives and market structure of political

manipulation techniques. Increasingly, state of the art social science involves methodological collaboration—computational analysis of data can reveal underlying patterns of information and behavior, but only in combination with ethnography can we undertake a theoretically meaningful interpretation of them. By combining both methods, this book considers both the broad patterns of computational propaganda in our online information environment and the deep-seated political, economic, and sociocultural forces that operate to encourage, promote, and leverage it.

Social Network Analysis. The research team applied social network analysis techniques to integrate large datasets from social networking platforms with existing survey and polling data, in order to understand how the structure of social networks bounds our opportunities for political learning, engaging with political leaders, and empathizing with social problems. Mapping advert networks on social media can reveal how far a misinformation campaign has traveled and what made it go viral, while mapping social networks can reveal how users may self-select into "echo chambers."

Surveys and Public Opinion Polling. Several of the chapters in this volume take advantage of existing survey instruments to trace the impact of algorithmic manipulation on public opinion. Surveys can be useful in generalizing user opinion on the political use of social media. Open-ended questionnaires allow users to elaborate on elements of digital political manipulation and harassment that researchers might not otherwise have encountered or considered. Scaled questions allow researchers to ascertain whether or not, and to what extent, users believe computational propaganda to be a problem they have experienced.

Comparative Policy Analysis. The legal research methodologies used in this edited collection examine how governments, and other jurisdictions where relevant, have implemented the key privacy and data protection principles and safeguards of leading privacy and data protection instruments. In most cases, national governments do not have electoral laws that help political actors make good decisions about the complex range of political communication technologies currently on offer. Wherever possible, chapter authors have discussed the rules that should apply in election campaigning—or that perhaps should have been applied.

Computational Propaganda: Addressing a Global Problem

We have already mentioned that the World Economic Forum has identified the rapid spread of misinformation online as one of the top 10 perils to society. In this book we present new, original evidence about how this manipulation and amplification of disinformation is produced, managed, and circulated by political operatives and governments. We measure how Russian Twitter conversation is constrained by highly automated accounts. We demonstrate how highly automated accounts in the United States have moved from peripheral social networks to engage with core groups of humans. We also trace the source of some forms of junk news and automated accounts to programmers and businesses in Germany, Poland, and the United States.

Our interviews with political party operatives, freelance campaigners, and elections officials in seven countries provide further evidence that social media bots—and computational propaganda more broadly—have been used to manipulate discussion online. This manipulation is underscored, and indeed facilitated, by the fact that some social media platforms, in particular political contexts, are either fully controlled by or dominated by governments and organized disinformation campaigns. Almost half of the Twitter activity in Russia is managed by highly automated accounts. Significant portions of political tweeting in Poland are produced by just a handful of alt-right accounts.

Computational propaganda also plays a role in particular events, especially during elections and security crises. It played a significant role during three recent political events in Brazil: the 2014 presidential elections, the impeachment of former president Dilma Rousseff, and the 2016 municipal elections in Rio de Janeiro. Analysis of how the Ukrainian conflict has played out on social media provides perhaps the most detailed case of computational propaganda's role during a global security crisis, and Russia's ongoing involvement in information wars. Numerous online disinformation campaigns have been waged against Ukrainian citizens on VKontakte, Facebook, and Twitter. The industry that drives these efforts at manipulation has been active in Ukraine since the early 2000s.

Computational propaganda also flourished during the 2016 US presidential election (Howard, Kollanyi, & Woolley, 2016), with numerous examples of misinformation distributed online with the intention of misleading voters or simply earning a profit. Multiple media reports have investigated how "fake news" may have propelled Donald J. Trump to victory (Dewey, 2016; Parkinson, 2016; Read, 2016). In Michigan, one of the key battleground states, junk news was shared just as widely as professional news in the days leading up to the election (Howard et al., 2017). Surveys have suggested that many people who saw fake news during the election believed those headlines (Silverman & Singer-Vine, 2016), though we have yet to see firm long-term interference with political learning.

There is a difference in how computational propaganda is used by authoritarian and democratic governments. Increasingly, however, this gap is closing. Our case studies show that authoritarian governments direct computational propaganda at both their own populations and at populations in other countries. Campaigns directed by China have targeted political actors in Taiwan, and Russian-directed campaigns have targeted political actors in Poland and Ukraine. But in democracies as well, individual users design and operate fake and highly automated social media accounts. Political candidates, campaigners, and lobbyists rent larger networks of social media accounts for purpose-built campaigns, while governments have assigned public resources to the creation, experimentation, and use of such accounts. And this doesn't just rely on automation and AI technology; when it comes to effective use of computational propaganda, the most powerful forms will involve both algorithmic distribution and human curation—software bots and human trolls working together. Our Taiwanese study reveals that Chinese mainland propaganda over social media is not fully automated but is in fact heavily coordinated by humans.

It's not all bad news. There are important examples of positive contributions made by algorithms and automation over social media. In Canada, civic actors are using complex algorithms to do constructive public service—albeit, with an as-yet uncertain overall impact. Bot builders in the United States have constructed small groupings of social bots with mandates for making information about political processes more public and easier to understand, for creating art aimed at critiquing particular policies or social issues, and for connecting social or political groups with similar interest groups.

Our motive in undertaking this multi-case analysis of computational propaganda is to better understand the global reach of political bots, digital disinformation, junk news, and other similar problems. In presenting the first systematic exposé and analysis of computational propaganda for a number of country-specific case studies, we have paid particular attention to the themes inherent in propaganda generally, but also try to illuminate crucial details surrounding particular attacks and events. Ultimately, we hope to understand who is behind misinformation campaigns while also explaining who the victim groups are, what they experience, and what they—and others fighting this global problem—can do about it.

Conclusion — Can Democracy Survive Computational Propaganda?

We find several distinct global trends in computational propaganda. While it is true that social media are significant platforms for political engagement, crucial channels for disseminating news content, and the primary media over which young people develop their political identities, they are also—and perhaps in part because of these affordances—vessels for control. In some countries this problem is exacerbated because companies such as Facebook have effectively become monopoly platforms for public life. In several democracies the majority of voters use social media to share political news and information, especially during elections (Bakshy, Messing, & Adamic, 2015). In countries where only small proportions of the public have regular access to social media, such platforms are still fundamental infrastructure for political conversation among the journalists, civil society leaders, and political elites (Farhi, 2009; Hermida, 2010). With this confluence of communication and sense making comes efforts to co-opt the flow of communication.

Social media are actively used as a tool for public opinion manipulation, though in diverse ways and on different topics. In authoritarian countries, social media platforms are a primary means of social control. This is especially true during political and security crises but is generally true in day to day life. In democracies, social media are actively used for computational propaganda either through broad efforts at opinion manipulation or targeted experiments on particular segments of the public. In every country we found civil society groups trying, but struggling, to protect themselves and respond to active misinformation campaigns.

We face new challenges in the investigation of automation and fake accounts on social media. Bots and sock-puppet accounts—fake accounts run by people—are key tools for spreading computational propaganda. Automation and anonymity allow for large scale amplification of some ideals or candidates for office alongside suppression of others. We have found that political actors are adapting their automation in response to our research. This suggests that the campaigners behind fake accounts and the people doing their “patriotic programming” are aware of the negative coverage that this gets in the news media.

We have also found several kinds of bot networks that are quite active but that fall below our formal threshold of what counts as a bot—or highly automated. For countries where Twitter is not a particularly important social media platform, it seems that bots are prevalent but not performing as efficiently as bot networks in countries with lots of Twitter users. Bots do not necessarily need to message at high rates in order to adversely affect public opinion or trending algorithms. Large numbers of automated accounts can be run by one person, converge on a topic or hashtag, and through this affect the flow of political communication.

Increasingly, bots supported via close attention from human operators—cyborg accounts—are being used to circumvent algorithms set to detect automation. Headless browsing bots get around these mechanisms by logging on to social media sites rather than via the application programming interface (API). Coordinated human-run accounts have also been successful in political hashtag bombing and trend manipulation (Musgrave, 2017). In many countries there are large numbers of “sleeper bots” (Woolley & Howard, 2016; Bradshaw & Howard, 2017). These are accounts that have only tweeted a few times, usually in scattered ways, and have other account features that suggest automation.

It is difficult to put research findings into service for public policy recommendations in consistent ways across countries, because the legal questions about computational propaganda vary greatly from country to country. During the 2015 election in Canada, comedienne Sarah Silverman encouraged Canadians to vote for the National Democratic Party over Twitter (Itzkoff, 2018). Is she a foreigner influencing voters in contravention of the Canada

Elections Act? If bots propagate her message after campaigning is supposed to stop, are platforms or bot writers interfering with the election? When political bots are built and launched using crowd-sourced or open-source code, who is responsible for their actions? Also, how can we preserve democratically beneficial bots? It has been argued that bots can act as social scaffolding for journalists and democratic activists (Hwang & Woolley, 2016; S. Woolley et al., 2016)? However, positive uses are threatened by attempts to prevent malicious uses of social automation.

The advantage of cross-national comparisons is in yielding evidence about which policy responses can work well. In Taiwan the government has responded with an aggressive media literacy campaign, and bots that will check facts for the public. In Ukraine the government response has been minimal, but there are a growing number of private firms trying to make a business of fact checking and protecting social media users—Youscan.io, ContextMedia, Noksfishes, SemanticForce, and InfoStream. In the United States, platforms like BotoMeter, NewsbotAI, and botcheck.me are becoming industry standards for detecting nefarious bots as well as disinformation. However, it is important to note, whether solutions are short term or long term, that global society needs fixes that are social as well as technological—the Taiwanese case being a good example of this two-pronged approach. Moreover, our research shows this is not simply an issue that can be solved by giving users access to more or better information. Companies and governments have a crucial role in combating computational propaganda through new policies and interventions.

Automated political communication involves the creation, transmission, and controlled mutation of significant political symbols over expansive social networks. Indeed, the impact of digital information infrastructure on how political culture is produced is at least as interesting, though under-studied, as the impact of infrastructure on how political culture is consumed. While we can theorize about the ways in which computational propaganda may violate political values or the social contract writ large, it is difficult to quantify these effects. But the case studies in this collection of working papers demonstrate the origins and very concrete consequences of computational propaganda.

It is time for social media firms to design for democracy. For democracies, there will always be big elections ahead. Let's assume that authoritarian governments will continue to use social media as a tool for political control. But for democracies, we should assume that encouraging people to vote is a good thing. Promoting political news and information from reputable outlets is crucial. Ultimately, designing for democracy, in systematic ways, will help restore trust in social media systems.

Computational propaganda is now one of the most powerful tools against democracy. Social media firms may not be creating this nasty content, but they are the platform for it. The Facebook Newsfeed, and Trending features on sites like Twitter and YouTube, produce curated content (Gillespie, 2010). This means that these features prioritize, or control, the information that people see. Because social media companies have made decisions to control what information or news people see, these entities have responsibility for making sure this information is not harmful, harassing, or false. This is especially true during pivotal political events like elections, but also true in general.

Platforms need to significantly redesign themselves if democracy is going to survive social media. Moreover, they cannot rely upon tired defenses about being technology not media companies. Trending features, algorithmic curation, and personalized news feeds mean that companies do, to use their language, arbitrate truth. Because they control information flow, they are media companies. To solve these problems, social media companies must confront their role as media platform. They must design for democracy.

Remaining Questions

Cross-country comparison is a powerful way of understanding real trends and lived political experience. Yet all of the case studies here have conclusions that beg more questions, some theoretical and others practical. How should democracies advance security while protecting privacy? When should reasonable forms of surveillance be implemented, and under what circumstances? What technological design principles might provide both collective and personal security? How can we make algorithmic decision-making transparent, fair, and accountable? How can we create regulations to keep pace with innovation to safeguard our rights to be treated fairly by algorithmic decision-making systems?"

Research Challenges Ahead

While the researchers in this collection have demonstrated the global spread of political propaganda that takes advantage of the affordances of social media algorithms, there are several important political communication research questions that need answering. We know very little about the actual influence of highly automated accounts on individual political attitudes, aspirations, and behaviors. In short, it is hard to demonstrate that any particular tweet, Facebook post, or other social media message has a direct effect on a voter. Notably, this test is one that many political communication researchers dismiss as misinformed when it comes to print, radio, or television, but it has reappeared as an expectation of social media research. But more broadly, making a causal claim from social media use to citizen engagement, trust in institutions, or voter sophistication is proving difficult to do even in countries for which there are significant amounts of data. In democracies across the global south, understanding these dynamics are an important research challenge.

Cross-Case Comparison

Table C.1 summarizes the various national contexts we have investigated, and the consequences of having important political actors in each country develop and apply algorithms and automation in public discourse. The precise applications also vary—from referenda and elections to policy debates and national security crises. There are also some disturbing similarities between the contexts and consequences among countries we would normally distinguish as being democracies or authoritarian regimes.

Table C.1 **Country Specific Breakdowns of Computational Propaganda**

Country	*Domestic Political Actors Involved*	*Foreign Political Actors Involved*	*Prominence of Computational Propaganda in Political Communication*	*Observations*
China	State	N/A	Low	Much more human-driven propaganda campaigns on behalf of the party than those facilitated by bots.

Taiwan	Parties	Chinese	High	Taiwan experiences propaganda over social media from the Chinese mainland, and likely the Chinese state. It is mostly human driven.
United States	Parties, Firms, Lobbyists, Civil Society Groups	Russian	Moderate	Computational propaganda played a significant role in the 2016 US presidential election, with numerous political actor types deploying and using bots in attempting to manipulate public opinion.
Ukraine	Parties, State, Firms, Civil Society Groups	Russian	High	Ukraine is on the frontline of computational propaganda, and myriad offline problems, from Russia. It is perhaps the most advanced, and worrisome, case of computational propaganda explored here.
Russia	Parties, State	N/A	High	The Russian government, and entities including the Internet Research Agency, have developed new strategies for computational propaganda and honed kompromat strategies from the Cold War in attempts to influence both domestic and foreign political events over social media.

Conclusion

Democracy itself is under assault from foreign governments and internal threats, such that democratic institutions may not flourish unless social data science puts our existing knowledge and theories about politics, public opinion, and political communication to work. These threats are current and urgent, and, if not understood and addressed in an agile manner, will further undermine European democracies. Given current trends, it is likely that some political actors will begin using machine learning applications to produce political content during elections, or fully fake videos that are indistinguishable from real news, to undermine the public confidence in shared information and reasoned debate. Most democratic governments are preparing their legal and regulatory responses. Yet, unintended consequences from over-regulation may be as damaging to democratic systems as the threats themselves.

Technology innovation often provides new opportunities to dream of possible political futures. It invariably inspires new research questions in those of us who study public life, especially when new information technologies appear to exacerbate social inequalities and cause social problems rather than mitigate or solve them. The causes and consequences of computational propaganda certainly vary from country to country, and we are eager to develop a large research community that is normatively committed to redress social inequalities, solve public problems, and strengthen democratic institutions.