

+++ *unpublished manuscript, please consult the authors before quoting* +++

## **Social bots in Germany's 2017 national election campaign: Theoretical, empirical and methodological implications**

Tobias R. Keller, U of Zurich, t.keller@ikmz.uzh.ch

Ulrike Klinger, FU Berlin, ulrike.klinger@fu-berlin.de

***Abstract:** Social bots mimic and potentially manipulate humans and their behaviours in social networks. The public sphere is especially vulnerable to their impacts. We first discuss their potential influence on the public sphere from a theoretical perspective. From an empirical perspective, we analysed all Twitter followers of seven German parties before ( $N = 1,180,362$ ) and during the 2017 electoral campaigns ( $N = 1,588,213$ ) regarding bot prevalence and activities. The results revealed that the share of social bots increased from 7.1% before to 9.9% during the election campaigns. The percentage of active social bots remained roughly the same. An analysis of the content distributed by both the most influential and the most active bots showed that they utilise few political hashtags, and that almost none referred to German politics. We discuss the results against the background of normative traditions of public sphere theories and address the methodological challenges bots pose in political communication.*

### **1. Introduction**

Social bots are computer programs that mimic and potentially manipulate humans and their behaviours in social networks (Ferrara, Varol, Davis, Menczer, & Flammini, 2016; Wagner, Mitter, Körner, & Strohmaier, 2012). Bots post in online forums and dating platforms, and 'like', comment, and share social media contributions. They are cheap tools to make content, topics, or actors appear more popular than they really are. They start and catalyse online phenomena to stir outrage and artificial hype, while neither people nor trending algorithms can discern them with full accuracy as being non-human. Social bots differ from more general bot software that deliver simple services around information retrieval, selection or the creation of personalised preferences without directly interacting with Internet users (Woolley, 2016). Bots present an opportunity to connect artificial intelligence with simple applications; there has been an exponential growth in the number of Twitter bots on the largest open-source

online code repository, GitHub – which enables people with few programming skills to deploy social bots for their (political) purposes (Kollanyi, 2016). Bots are active on many platforms; for Facebook Messenger alone, independent developers have created more than 100,000 bots, which perform a variety of services.<sup>1</sup>

Social bots are no longer a marginal phenomenon on social media platforms. On Twitter, 9% to 15% of users are estimated to be bots (Varol, Ferrara, Davis, Menczer, & Flammini, 2017). Their everyday occurrence makes analysis and reflection imperative from a social science perspective, and calls for a convergence of social and computational science approaches (Strohmaier & Zens, 2014). How do social bots influence and change social and political discourses that are invisible and indiscernible to Internet users? Do social bots endanger and challenge the interactive and participatory potential of digital communication in mass democracies? How should studies assess empirical evidence of social bots' prevalence, activities and impacts, and what does this mean for studies on political communication in social media generally?

From the perspective of political communication, we address social bots from three angles: theoretical, empirical and methodological. We begin with a review of the theory and current literature regarding social bots, to discuss how they work and why they may create problems for democratic processes and public communication, drawing on public sphere theories. We then present an empirical study of social bots among Twitter followers of Germany's political parties before and during the country's 2017 national election campaigns. From a methodological perspective, our paper raises the question how political communication scholars who work with social media data should deal with social bots.

---

<sup>1</sup> See <https://venturebeat.com/2017/04/18/facebook-messenger-hits-100000-bots/> and <https://messenger.fb.com> (last accessed on 1 November 2017)

## 2. Theory: Social bots, agency and normative models of the public sphere

From a theoretical perspective, social bots challenge many concepts that social scientists take for granted, for instance, the question what constitutes an *actor*. While some authors understand bots as “automated social actors” (Abokhodair, Yoo & McDonald, 2015, p. 2), the question of technology and non-human agency is theoretically more complex. Technologies designed and programmed by humans embody social values and business models; they are encoded with human intentions and have limited agency of their own. Their behaviour is human-like and human-guided, which makes them human-dependent rather than autonomous actors. For instance, people overcome their physical limitations by using bots to retweet messages under multiple personas. Social bots impact communal relationship types (*Vergemeinschaftung*), i.e. social relationships based on a sense of belonging together, as well as associative relationships (*Vergesellschaftung*), i.e. social relationships based on rational agreements (Weber, 1922, §9). How do we account for social structures and relations that include social machines? And, since we know that social bots generate a large percentage of Internet traffic (Zeifman, 2015) and interact with human users, what does this mean for the formation of social relations and society, especially the public sphere?

Bots have existed since the early days of computer science, particularly chatbots designed to interact with humans (Turing, 1950). They now have a myriad of possible functions determined by their creators. They may be programmed to assist (e.g. bots in depression treatment<sup>2</sup> or to edit typos on Wikipedia<sup>3</sup>), harmlessly collect data (e.g. Googlebot<sup>4</sup>) or

---

<sup>2</sup> Woebot for example, Fitzpatrick, Darcy, and Vierhile (2017)

<sup>3</sup> For an overview see the respective Wikipedia page <https://de.wikipedia.org/wiki/Wikipedia:Bots> (last accessed 31 January 2018)

trivially entertain (e.g. meme generator bots<sup>5</sup>). An increasing number of companies employ bots for customer care (Chakrabarti & Luger, 2015), or for quick dissemination of alerts (e.g. weather, traffic, police etc.) (see Guo & Chen, 2014). However, their activities in digital political communication are more ambivalent.

On the one hand, there are many useful tasks for social bots in political communication. Similar to bots helping people to choose new outfits, bots could help citizens identify their political preferences and match them with parties and candidates (e.g. in voting advice applications, such as Wahl-O-mat in Germany or Smartvote in Switzerland). On the other hand, problems start when bots operate in disguise, interacting with citizens, voters, and stakeholders without people knowing. Social bots can orchestrate campaigns to hype organisations, alter perceptions of political reality by spreading propaganda (Abokhodair et al., 2015; Boshmaf, Muslukhov, Beznosov & Ripeanu, 2011), disrupt government and organisational communication (Woolley & Howard, 2016), feign grassroots movements (e.g. Sri Lanka) (Rathnayake & Buente, 2017), and alter public opinion by simulating the popularity of or protest against topics or actors (Ferrara et al., 2016). This implies that an increasing amount of online communication is non-authentic, but at the same time intended to yield real consequences.

Social bots can have different functions based on the behaviours they are programmed for. They can be merely *passive*, connecting with a number of accounts in order to boost the number of followers and their interconnectedness without contributing content. In this way, they make actors appear more popular and socially acceptable than they really are, encouraging others to follow or ‘like’ them (the so-called bandwagon effect) (see Sundar,

---

<sup>4</sup> For example Googlebot <https://support.google.com/webmasters/answer/182072?hl=en> (last accessed 31 January 2018)

<sup>5</sup> For example the Meme Generator Bot on Facebook on <https://www.facebook.com/memegenBot/> (last accessed 31 January 2018)

2008; Sundar, Oeldorf-Hirsch, & Xu, 2008). Alternatively, they could be *active*, liking, sharing, retweeting, commenting and broadcasting information, interacting in debates and fuelling discussions.

Online public spheres, such as social media platforms and online forums, have become commonplace for deliberation, political discourse, and the articulation and aggregation of political interests. Whether or not participants realise it, their interactions are likely to be infiltrated by social bots and their agendas; as Mitter, Wagner, & Strohmaier (2013, p. 1) put it: “Without a deep understanding of the impact of such attacks, the potential of online social networks as an instrument for facilitating discourse or democratic processes is in jeopardy.” However, whether social bots pose a threat to the public sphere or democracy largely depends on the normative perspective. Table 1 describes Ferree, Gamson, Gerhards, & Rucht’s (2002) four models of the public sphere in modern democracies: representative-liberal, participatory-liberal, discursive and constructionist.

**Table 1:** Normative traditions of public spheres

	<b>Inclusion</b>	<b>Processes</b>	<b>Key bot problems</b>
<b>Representative-liberal</b>	Elites, Experts	Recurring exchange of political elites in elections	Quantitative misrepresentation of popularity
<b>Participatory-liberal</b>	Popular inclusion	Plural decision-making	Diffusion of fake political interests (astroturfing)
<b>Discursive</b>	Popular inclusion	Deliberation	Non-authentic, manufactured participants; Lack of mutual respect and rationality
<b>Constructionist</b>	Popular inclusion	Empowerment of marginalized actors, Expansion of political community	Non-authentic, manufactured participants

*Note.* Overview of the models of public sphere is based on Ferree et al. (2002: 316).

In a representative-liberal tradition, the public sphere is an elite-dominated, free and transparent forum that enables citizens to repeatedly choose (and replace) their representatives. In this instance, bots potentially disturb the key principles of proportionality

and transparency. By discreetly making some ideas and actors appear more popular than they really are, coverage of political actors becomes disproportional to their de facto citizen following. Although citizens may be unaware of the situation, it becomes impossible for them to take popularity cues (Porten-Cheé, Haßler, Jost, Eilders & Maurer Marcus, 2017) from political actors in social networks as a proxy of public opinion and their popularity among fellow citizens, and there may even be a conflict of popularity cues online (likes, retweets) and offline (polls, media coverage). Thus, in a representative-liberal perspective, bots are a problem when they distort political competition, intervene in campaigns, and influence elections' outcomes. If bots boost political parties or candidates' number of Twitter followers, Facebook friends or group members, they threaten the functioning of key democratic processes.

In a participatory-liberal tradition, the public sphere is a space for public discourse that seeks to achieve maximum popular inclusion – not only during election campaigns, but all the time. Voices should not be linked to proportionality, but to plurality: all interests and actors in a community should be included and heard. However, with increasing bot presence, the desired inclusion of grassroots movements may turn more to *astroturfing*, i.e. “grassroots support that is artificial because it is manufactured and does not arise spontaneously” (Klotz, 2007, p. 5). The principle of plurality is based on the premise of authentic interests and stakes in a society. Bots may insert non-authentic interests (interests no human or group in a society have ever voiced) and manipulated interests (fake interests that are manufactured to distort plurality). It becomes impossible for a society to monitor itself when machines disguised as societal members enter and manipulate the marketplace of ideas (e.g. Alexander, 2015). This means that bots are not only a problem because they lead to quantitative misrepresentations and make parties or candidates seem more popular than they are, but because they could

potentially give voice to non-existent ideas. In the functioning of public spheres, this is particularly relevant when bots send and multiply (i.e. retweet) political messages.

This aspect becomes even more toxic in the discursive tradition: “But when important normative questions are at stake, it is crucial that the discussion not be limited to actors at the centre of the political system. On such issues, a well-functioning public sphere should simultaneously include actors from the periphery as well [...]” (Ferree et al., 2002, p. 300). Habermas’s distinction between autonomous (*autochtone*) and power-regulated (*vermachtete*) actors from the periphery becomes obsolete when automated, manipulative and interest-driven bots enter a discourse. Bots have no intention to understand or consider others’ opinions, and their participation in political discourse only emphasises their creators’ lack of respect for deliberative processes. The idea that decisions are made collectively, and conflicts are resolved based on an argument’s quality rather than on the number of supporters for an argument, is irreconcilable with social bots. With bots, discourse becomes impossible; debate turns into a travesty.

Finally, a constructionist perspective on public spheres focuses not only on plurality and inclusion, but on difference and mutual recognition: “Recognition politics, sometimes called identity politics, creates a good public sphere by decentring dominant speakers and their assumptions of what is ‘natural’ ” (Ferree et al., 2002, p. 308). In this tradition, everything is political, whether it takes place in private or in public, i.e. wherever power structures appear. The normative objective is gives voice to the marginalised, contesting and breaking “the boundaries between the public and private” (p. 311). This notion of public conversation and democratic processes is particularly vulnerable to the participation of bots, because it seeks to empower previously silent voices and to include fringe groups and their political claims, and prefers narrative styles over rational, unemotional debate. Bots can boost popularity cues and

take on the identity of an assumed marginalised group. Constructionist visions of public discourse depend on authentic individuals contributing genuine perspectives from their life-worlds (*Lebenswelt*), which are easily infiltrated and undermined by bots.

Bots are not inherently evil forces, and they are not all problematic for the same reasons. Any assessment of their impacts must acknowledge their empirical behaviour patterns and a theoretical reflection on normative assumptions about political communication and the public sphere.

### **3. Literature review: What we know about social bots in political communication**

While these questions are being discussed in blogs and newspapers, communication research is only starting to focus on social bots. Initial social science research projects have made it clear that social bots are by no means a merely technical phenomenon, but change how Internet users interact and form social relations among each other, and with institutions, organisations and society. Social scientists have only recently begun to address their potential to intervene in election campaigns and to distort public communication and deliberation (Hegelich & Janetzko, 2016; Woolley, 2016).

Social bots exist to participate in human interaction and discourse, and are finding a fertile habitat in social media networks (Guilbeault, 2016). An increasing number of companies employ bots for customer care (Chakrabarti & Luger, 2015), but social bots can also manipulate political (and other) discourses. Approximately one-quarter of Donald Trump's Twitter followers during the 2016 U.S. presidential campaign were bots (Woolley & Howard, 2016b). By focusing on hashtags, Kollanyi, Howard and Woolley (2016) found that bots made up a large part of Twitter traffic during the campaign and privileged Trump messages over Clinton ones. Bessi and Ferrara (2016) found that social bots were present and did



influence the U.S. presidential campaign, with about 20% bots involved, generating about 20% of the political debate on Twitter. They also showed how easy it is to employ bots, even for unexperienced and non-tech-savvy users, since they are offered by diverse companies, sometimes even on a monthly subscription basis (p. 2). Bots intervened in the Brexit debate (Howard & Kollanyi, 2016), and the online petition for a second referendum on Brexit in June 2016 was “signed” by 77,000 bots (BBC, 2016). Bastos and Mercea (2017) discovered a network of 13,493 Twitterbots supporting the Leave EU campaign. Social bots drove the #MacronLeaks disinformation campaign: “the users who engaged with MacronLeaks are mostly foreigners with a pre-existing interest in alt-right topics and alternative news media, rather than French users with diverse political views. Concluding, anomalous account usage patterns suggest the possible existence of a black-market for reusable political disinformation bots” (Ferrara, 2017, p. 1). A study of Germany’s 2017 election campaigns at the Oxford Internet Institute found that “highly automated” tweeting increased from 5.7% to 7.4% between February and September 2017. It also compared data from other projects but with the same research design, finding between 5.2% and 16.5% automated tweeting in various campaigns (Neudert, Kollanyi & Howard, 2017). Previous studies outside the U.S. and European contexts found that right-wing parties and radical opposition parties used social bots more often than other parties, for instance in Japan (Schäfer et al., 2017) and Venezuela (Forelle, Howard, Monroy-Hernandez & Savage, 2015). Hegelich and Janetzko (2016) identified and analysed a botnet connected to the Ukraine conflict and showed that social bots have political agendas and act relatively autonomously on the basis of complex algorithms.

This is all the more relevant because experimental studies show that users perceive social bots as equally credible, competent, attractive and interactive to human agents (Edwards, Edwards, Spence & Shelton, 2014; Everett, Nurse & Erola, 2016). In this perspective, social bots can be

understood as new actors in digital political communication and a key element of what has been termed *computational propaganda*: “We define computational propaganda as the assemblage of social media platforms, autonomous agents, and big data tasked with the manipulation of public opinion” (Woolley & Howard, 2016a).

From previous empirical studies on social bots, we can conclude that bots are omnipresent on platforms, particularly on Twitter, and that they are being used to influence political and other debates.

This case study of how social bots interfere with the digital public sphere focusses on Germany’s 2017 national elections. German parties started their campaigning in the midst of the discussion about how social bots impacted the 2016 U.S. presidential election (e.g. Zhang, Wells, Wang & Rohe, 2017). We focus on the cross-sectional and longitudinal variation of bots among party followers:

*RQ1: How did the presence of bot-followers vary a) among the parties and b) before and during the election campaigns?*

Based on the findings of the aforementioned studies, we can assume that between 5% and 25% of Twitter accounts are bots (Bessi & Ferrara, 2016; Neudert et al., 2017; Varol et al., 2017). We can also assume that the number of bots is higher during a campaign phase than in a non-electoral period, especially since bots are sometimes removed from a platform after a campaign (Bastos & Mercea, 2017).

*H1: Of the Twitter accounts following German parties, 5% to 25% were social bots.*

*H2: The share of social bots was higher in election campaigns than in non-electoral periods.*

While passive social bots among a party's followers may increase their popularity, active bots' functionalities are more sophisticated. They are able to 'like' or retweet parties' messages, making them appear more popular and spreading them through the network. During election campaigns, the incentives to use bots to increase a party's visibility and to impact on political debate are higher, both for a party's supporters or other actors with an interest in influencing an election (e.g. Ferrara, 2017).

*H3: The share of active bots was higher during election campaigns than in non-electoral periods.*

All but one political party in Germany have pledged to not strategically use social bots during their campaigns, after social bots became a topic of public debate. One party, the right-wing populist AfD (Alternative for Germany) in October 2016 declared that "of course" it would implement social bots in its election campaign – "after all, for young parties such as ours social media tools are important instruments to proliferate our positions among voters"<sup>6</sup> (Stürzenhofecker, 2016, p. 1). A few days later, the party retracted this statement, declaring it would not to use bots in its campaign. However, since similar right-wing populist parties in Japan, Venezuela and France used bots (Ferrara, 2017; Forelle et al., 2015; Schäfer et al., 2017), and Neudert et al. (2017) found that most bots were supporting AfD, we hypothesise:

---

<sup>6</sup> Original quote: "Gerade für junge Parteien wie unsere sind Social-Media-Tools wichtige Instrumente, um unsere Positionen unter den Wählern zu verbreiten."

*H4: The right-wing populist party AfD had the largest share of social bots among its followers.*

Most studies measuring bots during election campaigns have remained silent on the content disseminated by bots (e.g. Zhang & Lu, 2016). Bots' presence alone is regarded as problematic. Nonetheless, it is possible that the same bots that boost parties and candidates' popularity levels spread no political messages, only commercial advertisements (Maireder, Weeks, Gil de Zúñiga, & Schlögl, 2016) or even nonsensical content (Bucher, 2017). We assume that bots post political content more often during an election period than in a non-electoral period.

*H5: Bots sought to influence Germany's political debate more during the election campaign by spreading political messages than in the non-electoral period.*

A final aspect in this literature review addresses the methodological approaches of bot-detection in studies. Bots avoid detection, and their creators invest effort into their resembling human users. Thus, it is not easy to identify bots and to distinguish these Twitter accounts from human accounts. Zhang and Lu's (2016) computational approach only used a user's network information to determine whether an account is a bot or a human on Weibo. Thus, they identified "millions of spammers" (Zhang & Lu, 2016, p. 14). The downside of their strategy is that they can only identify spambots depending on someone's network. Another approach, that of Hegelich and Janetzko (2016), is based on the URL that a tweet was sent from, which can be retrieved as part of a tweet's metadata. By manually identifying tweets sent from obvious bot creators, such as Twifarm, they searched for accounts that followed

these bots in order to unveil bot-networks. This approach is only possible if one starts with hashtags and tweets, not with Twitter accounts, because account metadata contains no URL information. The downside is that this method only detects a small amount of active bots. For instance, in Hegelich and Janetzko's case study, only 1,740 bots followed one another. Another approach is to identify and single out behavioural aspects that differ from human users, such as a high frequency of messages sent. Various studies on campaigns (e.g. Howard & Kollanyi, 2016; Neudert et al., 2017) came from a group of scholars who counted any account that sent out more than 50 tweets per day as bots, assuming "high automation". Needless to say, it is possible for a person to send 50 tweets per day (Musgrave, 2017). Also, not all bot accounts are this active, since passive bots only exist to boost certain accounts' follower numbers. Thus, this approach can only capture a specific bot type that broadcasts very actively. There is also the strategy of detecting social bots via near-duplicate tweets (i.e. copies or very similar versions of the same tweet sent by multiple bots), which bots use to inflate certain topics' frequency and importance (Schäfer, Evert & Heinrich, 2017). While this detection method is very useful, it remains unclear whether humans copied and pasted a tweet's content (e.g. Musgrave, 2017). Even if these copy-and-paste users were bots, they discover only one type of active copy-bots. Another share of studies uses multiple indicators to detect social bots: More elaborate ones use indicators such as "tweets to user", "mean tweet to retweet", "common words in the username" or "ratio of outbound to inbound @-mentions" (Bastos & Mercea, 2017, p. 6) to capture more than simple automated accounts. A similar approach, by Guo and Chen (2014), with a focus on geotagged tweets, propose four steps to identify spambots, including machine learning techniques. A drawback is the focus on geotagged tweets, because many Twitter users opt out of this option. Although very elaborate,

these techniques are hardly reproducible, since they require programming skills, which many social scientists (like us) lack.

With any bot-detection method, scholars face two key problems: 1) the cat-and-mouse game between bot creators and bot-detection developers and 2) the limited availability of data from commercial platforms. With complete datasets, it should be easy for platform owners to detect bots, but the incentive to do this and to delete these accounts is perhaps not worth pursuing: bot accounts do not buy anything and have no value for advertisers, but they keep the user numbers high. A sudden drop in platform users may unsettle shareholders. Thus, scholars must make the best of a tough situation. Here, we use the bot-detection tool Botometer, which was developed and maintained by computer scientists; it checks more than 1,000 variables of an account for features that are typical for bots (Davis, Varol, Ferrara, Flammini, & Menczer, 2016). This tool has been used in previous studies and is currently the most sophisticated, reliable and available instrument for bot-detection (see the next section on methods). Botometer is open for other scholars to detect social bots on Twitter so as to replicate our analysis.

#### **4. Sample and methods**

We collected data on all Twitter accounts that followed the five German parties represented in parliament: conservative CDU and CSU, social-democrat SPD, socialist Die Linke, and conservationists Die Gruenen. We also studied the liberal FDP, and right-wing populist AfD, which were considered likely to successfully (re-)enter Germany's Parliament in 2017 (and did). There were two data collection waves: the first, before campaigns started in January and February 2017, and the second during the week before Election Day: 24 September 2017.

For both waves, we first downloaded the Twitter account data of all followers of the seven German parties, including metadata such as their Twitter ID, screen name and numbers of followers, following and tweets (via BirdSong Analytics). Metadata also included information on account activity, i.e. whether or not a follower had been active in the past three months – that is, whether he or she tweeted, retweeted, liked or replied to a tweet. Downloading took place between 1 January and 13 February 2017 for the non-election period (N = 1.180.362), and from 12 to 14 September 2017 for the campaign period (N = 1.588.213).

In the second step, we used Botometer (formerly BotOrNot) to identify bots and distinguish them from humans in an automated analysis via Botometer’s API (Python 3.5). Since social bots constantly change their appearance, they are complicated to detect (Thieltges, Schmidt, & Hegelich, 2016). Botometer is a publicly available bot-detection instrument created and maintained by computer scientists at the University of Indiana. At the time of our study, it was the most sophisticated available instrument and has been used in several academic research projects, both by the creators and other scholars (e.g. Bessi & Ferrara, 2016; Ferrara, 2017). To ensure quality and comparability between the two waves, we kept in close contact with the computer scientists who maintain Botometer (and are greatly indebted to them for their kind support).

Botometer “generates more than 1,000 features using available metadata and information extracted from interaction patterns and content” (Davis et al., 2016, p. 2). These can be grouped into six main classes, all of which are vital in detecting social bots (Varol et al., 2017): *user* features include the number of followers and tweets produced by users; *friends* encompasses follower-friend relations such as retweeting and mentioning behaviours between one another; *network* characteristics include in-strength and out-strength (weighted degree) distributions, density and clustering; *content and language* features include statistics about

length and entropy of tweets and Part-of-Speech tagging; *sentiment* features encompass arousal, valence and dominance, happiness, polarisation, strength and emoticon scores (see Varol et al., 2017). From these, Botometer calculates a probability score between 0 (human) and 1 (bot) for each Twitter account. Overall, the tool has an accuracy of 0.86 and suggests that 9% to 15% of all Twitter users are bots (Varol et al., 2017). Because the tool is better equipped to identify humans than bots, our threshold for bots should be fairly high – not all accounts with a probability score over 0.5 should be counted as bots.

We conducted step three, data manipulation and cleansing, in R (R Core Team, 2017). Our final sample consisted of all follower accounts that Botometer could evaluate (see Tables 2 and 3). Three errors prevented us from calculating a final score for 541,688 follower accounts in the non-electoral and 838,026 accounts in the campaign phase: (1) an empty timeline (478,954 / 667,477 accounts), (2) the deletion of a Twitter account in the days between data collection and data analysis (1,269 / 2,135), and (3) privacy settings not authorising access to run an analysis (61,465 / 80,575). We were able to calculate a probability score for about 54% of followers in the non-electoral period, and 53% of followers in the campaign period (an interesting finding, which we discuss in the last section of the paper).

Table 2: Sample of the non-electoral period (January to February 2017)



	<b>Number of Followers (total)</b>	<b>Final Sample</b>	<b>Error (Sum)</b>	<i>Error: No Timeline</i>	<i>Error: Page does not exist anymore</i>	<i>Error: Not authorized</i>
AfD	47.534	<b>31.885</b>	15.649	11.946	12	3.691
CDU	161.025	<b>88.207</b>	72.818	61.766	796	10.256
CSU	123.324	<b>60.795</b>	62.529	56.607	36	5.886
FDP	148.311	<b>75.470</b>	72.841	66.119	33	6.689
GRUENE	290.679	<b>160.152</b>	130.527	115.769	144	14.614
LINKE	155.599	<b>86.447</b>	69.152	61.324	57	7.771
SPD	253.890	<b>135.718</b>	118.172	105.423	191	12.558
<b>Sum</b>	1.180.362	<b>638.674</b>	541.688	478.954	1.269	61.465

Table 3: Sample of the campaign period (September 2017)

	<b>Number of Followers (total)</b>	<b>Final Sample</b>	<b>Error (Sum)</b>	<i>Error: No Timeline</i>	<i>Error: Page does not exist anymore</i>	<i>Error: Not authorized</i>
AfD	74.923	<b>44.912</b>	30.011	24.541	228	5.242
CDU	221.114	<b>118.446</b>	102.668	89.469	306	12.893
CSU	166.631	<b>81.389</b>	85.242	76.875	221	8.146
FDP	244.624	<b>119.124</b>	125.500	114.340	337	10.823
GRUENE	356.481	<b>192.472</b>	164.009	146.104	358	17.547
LINKE	200.089	<b>110.578</b>	89.511	78.956	332	10.223
SPD	324.351	<b>171.105</b>	153.246	137.192	353	15.701
<b>Sum</b>	1.588.213	<b>838.026</b>	750.187	667.477	2135	80.575

To allow direct comparability between the different number of party followers and the metric scale scores, we calculated probability density functions (PDF), which we present as density plots to answer *RQ1*. To test *H1* to *H5*, we set the threshold for detecting bots depending on the density plots and accuracy of Botometer.

For a content analysis of information that bots distribute during election campaigns (*H5*), we selected the 100 most active bots (number of tweets sent) and the 100 most influential bots

(number of followers). These are considered the most interesting cases, since they could flood Twitter with political messages or function as opinion leaders. After having downloaded their tweets for the extended non-electoral period between 2 January and 2 April 2017 ( $N_{(\text{active})} = 60,262$ ,  $N_{(\text{influential})} = 42,425$ ), and the extended campaign period between 24 June and 24 September 2017 ( $N_{(\text{active})} = 36,804$ ,  $N_{(\text{influential})} = 14,130$ ), we extracted all hashtags used in these posts to assess the overall topics of these tweets (Bruns & Burgess, 2015; Small, 2011). The 100 most used hashtags for both the most active and most influential bots in both periods were manually examined by the authors. We analysed 400 hashtags to assess whether they were political (e.g. #Brexit), and checked against the tweets containing the hashtag for verification. This was particularly necessary for hashtags that function as non-political statements. For instance, tweets with the hashtag #Quotes focus mainly on various non-political quotes, seldom cite a popular statement by a politician. We did not count such hashtags as political, since they weren't political. The pre-test of 50 hashtags showed good reliability of the coders with a Krippendorff's alpha of 0.87.

## 5. Results


Regarding *RQ1*, we compared the distribution of scores between 0 (human) and 1 (bot) across the seven German parties' followers. We found three different patterns in the non-electoral period; while in the campaign period, all parties except the right-wing populist AfD show a similar pattern (see Figure 1 and Figure 2).

During the non-electoral period, CDU, CSU and FDP had a similar score distribution, as did SPD, Linke and Gruene; AfD showed a unique pattern. In the election phase, the general score distribution moved to the right, indicating a tendency of more bots among party followers. This is also true for AfD, whose follower distribution changed only marginally

between the non-electoral and the campaign periods, although they gained more followers in the time between the two waves (+27,389 followers, +57%).


The strongest differences could be found regarding accounts that are most likely bots, i.e. followers with scores above 0.75. In the non-electoral period, SPD, Linke and Gruene showed a peak around 0.75, all other parties expressed a smaller peak between 0.85 and 1. During the campaign period, this peak between 0.85 and 1 was more pronounced for all parties. Notably, the probability density function pattern during the campaign period was very similar for all parties, except AfD. One reason for this, we discovered, is that most parties share followers. While AfD had about 45% single-party followers (i.e. followers that only follow AfD but no other party), the average share of single-party followers among all other parties was 17%. This means that approximately 83% of Twitter users in our sample (whether human or bot) followed multiple parties – not exactly a sign of echo chambers. The share of single-party followers decreased slightly from February to September (-2%), i.e. during the campaigns, more followers chose to also follow other parties. This also indicates that a surplus of followers does not necessarily mean that more people follow parties on Twitter, but that people who already follow a party also follow other parties. Among the 407,851 new followers the parties acquired between February and September, only 79,272 (19%) were single-party followers.

**Figure 1:** All followers' probability of being a bot in the non-election phase in a density plot  
(January and February 2017)



*Note.* Density plot of all seven German parties' followers and their Botometer scores in the non-election phase.  $N = 638.674$ , bandwidth = ndr0 (see Silverman, 1986). The area between two scores and the function = the probability of a follower receiving such a score.

**Figure 2:** All followers' probability of being a bot in the campaign period in a density plot (September 2017)



*Note.* Probability density function (PDF) of all seven parties' followers and their Botometer scores in the election phase.  $N = 838.026$ , bandwidth = ndr0 (see Silverman, 1986). The area between two scores and the function = the probability of a follower receiving such a score.

Considering Botometer's accuracy, the hard task to distinguish between humans and bots in general, and the distribution of the scores in our two waves, we set the threshold for bots at a

score of 0.76. We found that the share of social bots fall mostly in the expected range of 5% to 25%, confirming *H1* (see Table 4).

Comparing the two waves, almost all parties had more followers, and most parties had more social bots among their followers in the campaign period than in the non-electoral period: AfD gained 0.4% more bots, Gruene 4.2%, Die Linke 8.3% and SPD 8.2%. Some parties had fewer bots during the campaign period: CDU -0.3% bots, CSU -0.8% and FDP -0.2%. Overall, the mean share of social bots among the seven German parties' followers rose from 7.1% to 9.9% (11,105 to 23,373 social bots), generally supporting *H2*, but not for all parties. Of these social bots, 212 (2%) were active during the non-electoral period and 314 (1.4%) during the campaign period. The numbers for each party are reported in Table 4. Social bots were not more active during the campaign period, leading us to reject *H3*.

Regarding *H4*, that the populist party AfD had the most social bots among its followers, our analysis revealed that the opposite is true: During the campaign, AfD had the smallest share of bots among its followers (7.1%, 5,325 social bots). In the non-electoral period, AfD had a below-average share of social bots (6.7%, 3,181); however, AfD bots were particularly active: AfD had the second highest share of active social bots, with 2.8% during the non-electoral period and 1.7% during the campaign period. The data suggests that AfD had a larger share of human scores than all the other parties (scores lower than 0.4), a much higher share of active followers (33%/26% compared to an average of 15%/14.5% for the other parties), and the highest share of active humans. Thus, we had to reject *H4*.

**Table 4.** Bot presence and activity on Twitter in the non-electoral and campaign periods of Germany's 2017 national election

	Non-Electoral Period		Campaign period	
	Social bots among followers (%)	Active bots among their social bots (%)	Social bots among followers (%)	Active bots among their social bots (%)
AfD	3.181 (6,7 %)	88 (2,8 %)	5.325 (7,1 %)	90 (1,7 %)
CDU	16.419 (10,2 %)	399 (2,4 %)	21.981 (9,9 %)	363 (1,7 %)
CSU	13.759 (11,2 %)	253 (1,8 %)	17.238 (10,4 %)	236 (1,4%)
FDP	16.180 (10,9 %)	287 (1,8 %)	26.087 (10,7 %)	338 (1,3 %)
GRU	19.287 (6,6 %)	196 (1,0 %)	38.549 (10,8 %)	485 (1,3 %)
LIN	1.696 (1,1 %)	9 (0,5 %)	18.734 (9,4 %)	243 (1,3 %)
SPD	7.212 (2,8 %)	267 (3,7 %)	35.697 (11 %)	445 (1,3 %)
<b>Mean</b>	<b>11.105</b> <b>(7,1 %)</b>	<b>212</b> <b>(2 %)</b>	<b>23.373</b> <b>(9,9 %)</b>	<b>314</b> <b>(1,4 %)</b>

Finally, we analysed the content of the tweets from social bots in the non-electoral and the campaign periods. Of the 100 most frequent hashtags distributed by social bots following a German party in the non-electoral period, the influential bots used 13 political hashtags (the frequency of a hashtag's use ranged from 34 to 2,375, median = 71.5) and the most active bots 30 (the frequency ranged from 30 to 1965, median = 65.5). With one exception, these political hashtags did not refer to German politics. They covered issues concerning politics in Austria, the EU (without a focus on Germany), France, Great Britain, Nigeria and the U.S. The one exception was one very influential bot that distributed the hashtag #AfD to promote its political agenda, with a total of 53 tweets. Most non-political hashtags and tweets were adverts (e.g. ads for jobs, paintings, financial investments, etc.).

During the campaign period, the use of political hashtags decreased. Of the 100 hashtags analysed for each set, the most influential bots used only seven political hashtags (the frequency ranged from 12 to 2,263, median = 23), the most active bots only eight (the frequency ranged from 42 to 1,916, median = 73.5). Again, none of the political hashtags referred to German politics; most concerned U.S. finance and climate change politics (e.g. #MisesInstituteUSA, #DavidStockmansContraCorner, #environment, #green, #climatechange). Fewer hashtags related to Nigerian politics (e.g. #Biafra). Similar to the non-electoral period, most hashtags had a promotional purpose (e.g. #yoga, #realestate, #porn, #software, etc.) in the campaign period.

These results led us to reject *H5*. Neither the most influential bots nor the most active bots tweeted more political content during the election campaign than during the non-electoral period. In fact, only one hashtag (#AfD) in the non-electoral period pointed to bot-influenced content in Germany's politics.

## **6. Discussion**

Social bots in the digital public sphere pose at least three challenges for political communication research: theoretical challenges to established concepts of social science, empirical challenges of detection and the measurement of impacts, and methodological challenges to the general validity of popularity cues and social media analysis.

In summary, we analysed all Twitter follower accounts of seven German parties before and during the 2017 electoral campaign. The analysis confirmed two hypotheses, showing that the share of social bots among these parties' Twitter followers increased from 7.1% before to 9.9% during the election campaigns. Three hypotheses were rejected: the share of active



social bots did not increase during the election campaigns; AfD did not have more bot-followers than the other parties – on the contrary, it had the smallest share. The bots that we did identify distributed almost no hashtags connected to German politics. We will now discuss some implications of these findings.

Connecting our findings to Ferree et al.'s (2002) four normative models of the public sphere, the potential damage caused by social bots in election campaigns covers a spectrum of problems. From a representative-liberal perspective, the results show that bots caused a quantitative misrepresentation of popularity, because roughly 10% of Twitter followers were bots disguised as humans. In particular, in the case of SPD and Linke, the share of bots among followers increased by more than 8% during the campaigns. Thus, their follower growth of 22% each appears bigger than it really was if we only count humans as authentic followers. Thus, social bots did manipulate popularity cues, disturbing the principles of proportionality and transparency during the campaigns. However, their impacts remains purely in numbers, because there was hardly any political content spread by bots that related to the election. From a participatory-liberal and discursive perspective, it is interesting that the share of bots increased during election campaigns. In this tradition, the focus is much less on elections than on popular inclusion and authentic debate with genuine contributions at all times. The share of active bots (i.e. bots that like, share, comment and discuss) was very low: 2% and less in both waves. While the proponents of participatory-liberal and discursive understandings of public spheres would not exculpate bots as non-authentic participants in political debate, the low bot activity and their predominant distribution of non-political content would certainly be a consolation for them. From a constructionist perspective, the low bot activity is a greater reason for concern. Seeing it as crucial to include previously silent, marginalised voices in public discourse, the presence of software actors deliberately designed to manipulate

popularity cues or contributions totally undermines the notion of a public sphere, whether during an election campaign or at any other time, whether or not they are active. In sum, hardly any harm was caused by social bots in Germany's 2017 national election campaigns. However, bots were clearly present, and their omnipresence on social media platforms, combined with their role in other campaigns, should keep social media researchers alert, providing a sound reason to closely monitor their activities.

Methodologically, social bots challenge the validity of social media studies: if a large part of likes, tweets, shares and comments originates from bots, how can results from quantitative studies measuring political actors' interactivity and popularity on social media be validated? The findings show that even the increase in followers needs further differentiation: Are new followers really 'new'? Are they persons? We propose that a standardised test for bot activity should become part of empirical studies about political communication on social media, in order to ensure results' validity. To address this challenge, social scientists must cooperate with computer scientists and push for more and better tools to monitor bot activity on social media platforms. This also entails a more critical stance towards the validity of data from social networks – a key question for research quality, and not only in political communication. When large numbers of interactions on social media platforms are generated by bots, this must be reflected in the results and conclusions of studies based on this data type, for instance, network analyses of political actors or analyses of campaign communication on social media platforms.

As with all single-case studies, this analysis has limitations. We examined only one election campaign, in one country, on one platform. Future studies should compare multiple countries and different media as well as political systems. In a connected world, foreign actors may try to influence elections via social bots – they may already have done so, and successfully – so

that metadata on geographic locations could be included in the analysis. One could also start with hashtags instead of follower accounts, because tweets provide better geodata than accounts. Further, Botometer's bot-detection is not 100% accurate. Building social bots that mimic human behaviour and building tools to identify them is a cat-and-mouse game. Although Botometer is currently the most sophisticated public tool to identify social bots, there is uncertainty about the accurate identification of bots; followers with a score around 0.6 remain hard to classify. The grey area is even larger when including Twitter accounts that produced an error: almost 89% of them had never sent a tweet ("no timeline error"). Is the sole purpose of these dead accounts or bots to make parties appear more popular than they really are? How many of the private accounts, which we were not authorised to analyse, were bots actively tweeting behind the curtain of privacy settings?

Clearly, we need more empirical research into bots, their activities and impacts. Future studies should ask how Internet users make sense of and construct perceptions of reality from their online interactions. How much do they know about social bots' presence and activity? Are they aware that they are interacting with social bots? Do they recall instances of interaction with bots? What are their perceptions of and opinions about social bots? Because people know what they know and what to think about from mass media (Luhmann, 2000), scholars should also analyse how journalists cover social bots. Are bots on media agendas? How is coverage about bots framed (e.g. as a technological phenomenon or a social issue)? Finally, researchers must remain alert to social bots and their influences on established theoretical concepts, digital empirical data and current methods when analysing digital communication.

## References

Abokhodair, N., Yoo, D., & McDonald, D. W. (2015). Dissecting a Social Botnet: Growth, Content and Influence in Twitter. In *CSCW '15 Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing* (pp. 839–851). <https://doi.org/10.1145/2675133.2675208>

- Alexander, L. (2015, April 2). Social Network Analysis Reveals Full Scale of Kremlin's Twitter Bot Campaign. <https://globalvoices.org/2015/04/02/analyzing-kremlin-twitter-bots/>
- Bastos, M. T., & Mercea, D. (2017). The Brexit Botnet and User-Generated Hyperpartisan News. *Social Science Computer Review*, 0. <https://doi.org/10.1177/0894439317734157>
- BBC. (2016). *EU referendum petition hijacked by bots*. Retrieved from <http://www.bbc.com/news/technology-36640459>
- Bessi, A., & Ferrara, E. (2016). Social Bots Distort the 2016 US Presidential Election Online Discussion. *First Monday*, 21(11). Retrieved from <https://ssrn.com/abstract=2982233>
- Boshmaf, Y., Muslukhov, I., Beznosov, K., & Ripeanu, M. (2011). The Socialbot Network: When Bots Socialize for Fame and Money. In R. H. Zakon (Ed.), *Proceedings of the 27th Annual Computer Security Applications Conference*. New York, NY: ACM. Retrieved from <http://lersse-dl.ece.ubc.ca/record/264>
- Bruns, A., & Burgess, J. (2015). Twitter Hashtags from Ad Hoc to Calculated Publics. In N. Rambukkana (Ed.), *Digital Formations: v.103. HashtagPublics: The Power and Politics of Discursive Networks* (pp. 13–28). New York, NY: Peter Lang Publishing Inc.
- Bucher, T. (2017). About a Bot: Hoax, Fake, Performance Art. *M/C Journal*, 17(3). Retrieved from <http://journal.media-culture.org.au/index.php/mcjournal/article/view/814>
- Chakrabarti, C., & Luger, G. F. (2015). Artificial conversations for customer service chatter bots: Architecture, algorithms, and evaluation metrics. *Expert Systems with Applications*, 42(20), 6878–6897. <https://doi.org/10.1016/j.eswa.2015.04.067>
- Davis, C. A., Varol, O., Ferrara, E., Flammini, A., & Menczer, F. (2016). BotOrNot: A System to Evaluate Social Bots. In J. Bourdeau, J. A. Hendler, R. N. Nkambou, I. Horrocks, & B. Y. Zhao (Eds.), *The 25th International Conference Companion* (pp. 273–274). <https://doi.org/10.1145/2872518.2889302>
- Edwards, C., Edwards, A., Spence, P. R., & Shelton, A. K. (2014). Is that a bot running the social media feed? Testing the differences in perceptions of communication quality for a human agent and a bot agent on Twitter. *Computers in Human Behavior*, 33, 372–376. <https://doi.org/10.1016/j.chb.2013.08.013>
- Everett, R. M., Nurse, J. R. C., & Erola, A. (2016). The anatomy of online deception. In S. Ossowski (Ed.), *The 31st Annual ACM Symposium* (pp. 1115–1120). <https://doi.org/10.1145/2851613.2851813>
- Ferrara, E. (2017). Disinformation and social bot operations in the run up to the 2017 French presidential election. *First Monday*, 22(8). <https://doi.org/10.5210/fm.v22i8.8005>
- Ferrara, E., Varol, O., Davis, C., Menczer, F., & Flammini, A. (2016). The rise of social bots. *Communications of the ACM*, 59(7), 96–104. <https://doi.org/10.1145/2818717>
- Ferree, M., Gamson, W., Gerhards, J., & Rucht, D. (2002). Four Models of the Public Sphere in Modern Democracies. *Theory and Society*, 31(3), 289–324. Retrieved from <http://www.jstor.org/stable/658129>
- Forelle, M. C., Howard, P. N., Monroy-Hernandez, A., & Savage, S. (2015). Political Bots and the Manipulation of Public Opinion in Venezuela. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.2635800>
- Guilbeault, D. (2016). Automation, Algorithms, and Politics| Growing Bot Security: An Ecological View of Bot Agency. *International Journal of Communication*, 10, 5003–5021. Retrieved from <http://ijoc.org/index.php/ijoc/article/view/6135>
- Guo, D., & Chen, C. (2014). Detecting Non-personal and Spam Users on Geo-tagged Twitter Network. *Transactions in GIS*, 18(3), 370–384. <https://doi.org/10.1111/tgis.12101>
- Hegelich, S., & Janetzko, D. (2016). Are Social Bots on Twitter Political Actors? Empirical Evidence from a Ukrainian Social Botnet. In *Proceedings of the Tenth International AAAI Conference on Web and Social Media (ICWSM 2016)* (pp. 579–582).
- Howard, P. N., & Kollanyi, B. (2016). Bots, #StrongerIn, and #Brexit: Computational Propaganda during the UK-EU Referendum. *SSRN Electronic Journal*. <http://doi.org/10.2139/ssrn.2798311>
- Klotz, R. J. (2007). Internet Campaigning for Grassroots and Astroturf Support. *Social Science Computer Review*, 25(1), 3–12. <https://doi.org/10.1177/0894439306289105>
- Kollanyi, B. (2016). Where Do Bots Come From? An Analysis of Bot Codes Shared on GitHub. *International Journal of Communication*, 10, 4932–4951.
- Kollanyi, B., Howard, P. N., & Woolley, S. C. (2016). Bots and Automation over Twitter during the U.S. Election. Retrieved from <http://comprop.oii.ox.ac.uk/wp-content/uploads/sites/89/2016/11/Data-Memo-US-Election.pdf>
- Luhmann, N. (2000). *The reality of the mass media. Cultural memory in the present*. Stanford, CA: Stanford University Press.
- Maireder, A., Weeks, B. E., Gil de Zúñiga, H., & Schlögl, S. (2016). Big Data and Political Social Networks. *Social Science Computer Review*, 35(1), 126–141. <https://doi.org/10.1177/0894439315617262>

- Mitter, S., Wagner, C., & Strohmaier, M. (2013). Understanding The Impact Of Socialbot Attacks In Online Social Networks. In H. Davis (Ed.), *Proceedings of the 5th Annual ACM Web Science Conference* (pp. 1–6). New York, NY: ACM.
- Musgrave, S. (2017, August 9). 'I Get Called a Russian Bot 50 Times a Day'. *Politico*. Retrieved from <https://www.politico.com/magazine/story/2017/08/09/twitter-trump-train-maga-echo-chamber-215470>
- Neudert, L.-M., Kollanyi, B., & Howard, P. N. (2017). Junk News and Bots during the German Parliamentary Election: What are German Voters Sharing over Twitter? Retrieved from [http://comprop.oii.ox.ac.uk/wp-content/uploads/sites/89/2017/09/ComProp\\_GermanElections\\_Sep2017v5.pdf](http://comprop.oii.ox.ac.uk/wp-content/uploads/sites/89/2017/09/ComProp_GermanElections_Sep2017v5.pdf)
- Porten-Che  , P., Ha  ler, J., Jost, P. B., Eilders, C., & Maurer Marcus. (2017). Popularity Cues in Online Media: Theoretical and Methodological Perspectives in Political Communication Research. *Studies in Communication | Media*, (accepted for publication).
- R Core Team. (2017). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. Retrieved from <http://www.R-project.org/>
- Rathnayake, C., & Bunte, W. (2017). Incidental Effects of Automated Retweeting: An Exploratory Network Perspective on Bot Activity During Sri Lanka's Presidential Election in 2015. *Bulletin of Science, Technology & Society*, 37(1), 57–65. <https://doi.org/10.1177/0270467617736219>
- Sch  fer, F., Evert, S., & Heinrich, P. (2017). Japan's 2014 General Election: Political Bots, Right-Wing Internet Activism, and Prime Minister Shinz   Abe's Hidden Nationalist Agenda. *Big data*, 5(4), 294–309. <https://doi.org/10.1089/big.2017.0049>
- Silverman, B. W. (1986). *Density estimation for statistics and data analysis. Monographs on statistics and applied probability: Vol. 26*. Boca Raton, FL: Chapman & Hall/CRC.
- Small, T. A. (2011). What the hashtag? *Information, Communication & Society*, 14(6), 872–895. <https://doi.org/10.1080/1369118X.2011.554572>
- Strohmaier, M., & Zens, M. (2014). Analyse Sozialer Medien an der Schnittstelle zwischen Informatik und Sozialwissenschaften. In C. K  nig, M. Stahl & E. Wiegand (Eds.), *Soziale Medien* (pp. 73–95). Wiesbaden: Springer Fachmedien Wiesbaden. [https://doi.org/10.1007/978-3-658-05327-7\\_4](https://doi.org/10.1007/978-3-658-05327-7_4)
- St  rzenhofecker, M. (2016, October 21). AfD will Social Bots im Wahlkampf einsetzen. *Zeit Online*. Retrieved from <http://www.zeit.de/digital/internet/2016-10/bundestagswahlkampf-2017-afd-social-bots>
- Sundar, S. S. (2008). The MAIN Model: A heuristic Approach to Understanding Technology Effects on Credibility. In M. J. Metzger & A. J. Flanagin (Eds.), *The John D. and Catherine T. Macarthur Foundation series on digital media and learning. Digital media, youth, and credibility* (pp. 73–100). Cambridge, MA: MIT Press. <https://doi.org/10.1162/dmal.9780262562324.073>
- Sundar, S. S., Oeldorf-Hirsch, A., & Xu, Q. (2008). The bandwagon effect of collaborative filtering technology. In M. Czerwinski, A. Lund & D. Tan (Eds.), *Proceeding of the twenty-sixth annual CHI conference extended abstracts on Human factors in computing systems - CHI '08* (p. 3453). New York, NY: ACM Press. <https://doi.org/10.1145/1358628.1358873>
- Thieltges, A., Schmidt, F., & Hegelich, S. (2016). The Devil's Triangle: Ethical Considerations on Developing Bot Detection Methods. In *AAAI Spring Symposium Series* (pp. 254–257).
- Turing, A. M. (1950). Computing Machinery and Intelligence. *Mind*, 49, 433–460. Retrieved from <https://www.csee.umbc.edu/courses/471/papers/turing.pdf>
- Varol, O., Ferrara, E., Davis, C. A., Menczer, F., & Flammini. (2017). Online Human-Bot Interactions: Detection, Estimation, and Characterization. In *Proceedings of the Eleventh International AAAI Conference on Web and Social Media (ICWSM 2017)* (pp. 280–289). Retrieved from <https://aaai.org/ocs/index.php/ICWSM/ICWSM17/paper/view/15587/14817>
- Wagner, C., Mitter, S., K  rner, C., & Strohmaier, M. (2012). When social bots attack: Modeling susceptibility of users in online social networks. In *#MSM2012 Workshop proceedings* (pp. 41–48). Retrieved from [http://ceur-ws.org/Vol-838/paper\\_11.pdf](http://ceur-ws.org/Vol-838/paper_11.pdf)
- Weber, M. (1922). *Wirtschaft und Gesellschaft*. T  bingen, Germany: J.C.B Mohr (Paul Siebeck). Retrieved from <http://www.unilibRARY.com/ebooks/Weber%2C%20Max%20-%20Wirtschaft%20und%20Gesellschaft.pdf>
- 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/6161/5300>
- Woolley, S. C., & Howard, P. N. (2016a). Automation, Algorithms, and Politics| Political Communication, Computational Propaganda, and Autonomous Agents — Introduction. *International Journal of Communication*, 10, 4882–4890. Retrieved from <http://ijoc.org/index.php/ijoc/article/view/6298/1809>
- Woolley, S. C., & Howard, P. N. (2016b, May 15). Bots unite to automate the presidential election. *WIRED*. Retrieved from <https://www.wired.com/2016/05/twitterbots-2/>
- Zeifman, I. (2015). 2015 Bot traffic report: Humans take back the web, bad bots not giving any ground. Retrieved from <https://www.incapsula.com/blog/bot-traffic-report-2015.html>

- Zhang, Y., & Lu, J. (2016). Discover millions of fake followers in Weibo. *Social Network Analysis and Mining*, 6(1), 1–16. <https://doi.org/10.1007/s13278-016-0324-2>
- Zhang, Y., Wells, C., Wang, S., & Rohe, K. (2017). Attention and amplification in the hybrid media system: The composition and activity of Donald Trump's Twitter following during the 2016 presidential election. *New Media & Society*, 1(1). <https://doi.org/10.1177/1461444817744390>