

Debate

Tommy Hasert / Gabriele Hooffacker

Overrated bots?

An examination of Twitter debates – and what journalists can learn from it

Abstract: Social bots are suspected of having an impact on public discourse, manipulating election results, and seeking to influence political conflicts. This paper is based on an investigation that sought to detect and evaluate social bots in current Twitter debates. The authors show that the influence of bots appears much less dramatic than is often written about. In fact, over-regulation presents a greater threat to democracy than the bots themselves.

In participative platforms with user-generated content, such as Twitter or Facebook, it is not always immediately obvious whether one is interacting with a real person or with an algorithm-controlled account that imitates human activities – known as a social bot.

Bots are said to have a huge influence. British newspaper *The Guardian* once ran the headline: »Social media bots threaten democracy« (Woolley/Gorbis 2017). In an opinion piece, the *Washington Post* even questioned whether the American democracy could survive this kind of interference in social media: »Artificial intelligence is transforming social media. Can American democracy survive?« (Watts 2018). »Social Bots – eine Gefahr für die Demokratie [A danger to democracy],« wrote Christian Kerl in the *Berliner Morgenpost* (Kerl 2017). Other German media argue that social bots have enormous destructive potential. On October 24, 2016, for example, *Spiegel Online* ran the headline: »Social Bots – Wie digitale Dreckschleudern Meinung machen [How digital muck-spreaders shape opinion]« (Amann et al. 2016). According to them, events in world politics such as Brexit, the 2016 US presidential election campaign, and the Russia-Ukraine conflict were influenced by social bots.

The attention this gained in society and the media led to calls from politicians for greater regulation of social bots. Proposals ranged from a simple labelling obligation (from political party Die Grünen) to a general obligation to use real names online (CDU/CSU). Politicians claimed that bots had the potential to curtail fundamental freedoms for all internet users (cf. Reuter 2017, 2019).

However, only a tiny part of this discourse in politics and the media is based on empirically substantiated figures. This paper therefore examines how widespread social bots really are on Twitter (Facebook was not included in the investigation) and what influence their activities have.

Bot strategies

How do bots influence social actors, opinions, and debates? There are four key strategies:

Imitation

Imitating human behavior is a fundamental property that defines every bot. The primary goal is to trick people into trusting the computer-controlled account by simulating a human identity. Social bots do this by basing themselves on human behavior in social networks. On Twitter, for example, it means tweeting and retweeting, following other users, participating in discussions by replying, and adding external content to favorites (cf. Misener 2011).

Language also plays a vital role. Using natural language algorithms is intended to produce more authentic responses (cf. Good 2016). Spelling mistakes and slang are deliberately incorporated to make the bots more difficult to uncover: »To avoid detection, they may even employ slang words, street idiom and socially accepted spelling mistakes« (Woolley/Howard 2014).

The use of persuasive techniques also seems to be a popular method of increasing trust in the bot accounts. Given the social proof effect, for example, it is a good idea to gain as many followers as possible at the start, as these act as social proof of the relevance of the account and thus increase its influence. A study by Hegelich and Janetzko showed that most of the content disseminated by a network of social bots investigated did not pursue a direct mission, but was solely intended to gain trust and maintain the cover (cf. Hegelich/Janetzko 2016). Once this trust/influence has been developed, it can be used to advance individual objectives, such as promoting a particular political opinion or discrediting a public figure.

Simulating trends

Social networks increasingly act as trend barometers for the relevance of topics in society. The last decade has seen the emergence of an entire industry dedicated to identifying trends on social media so that they can be exploited for commercial gain. Journalists and politicians also use the various social media channels to gain a sense of »what makes sections of society tick« (Weck 2016). Social bots benefit from this social relevance and generate their own trends through the intense use of certain hashtags or key words. This can create a distorted picture of the true real significance of a topic (cf. Meiselwitz 2017: 381).

Astroturfing

Astroturfing describes the strategy of »organizing particular interests as supposed desires of citizens [...] with the aim of influencing sociopolitical decisions« (Irmisch 2011: 24). Social bots use this approach within public debates in order to create the impression that there is a great deal of support for certain opinions or movements, when this is not actually the case. This can lead to strategic distortion of a debate. The structure of social networks like Twitter and Facebook, which offer the option of sharing and liking other content as a core function, provides the ideal conditions for this. A particularly popular strategy for gaining influence in political contexts, its existence has already been proven numerous times in various studies (cf. Ratkiewicz et al. 2011).

Smoke screening

Smoke screening is another influencing strategy used by social bots. Unwelcome debates are disrupted in a targeted way through dissemination of huge numbers of discrediting or irrelevant messages. In the context of Twitter, this approach is also referred to as a »Twitter bomb« (cf. Brachten et al. 2017), as it makes individual hashtags unusable. A study by Abokhodair et al. into a Syrian bot network was able to prove the use of smoke screening (cf. Abokhodair et al. 2015).

How can bots be uncovered?

As well as basic obstacles like user authentication and crowd-based approaches (reporting system), the large social networks predominantly rely on bot detection processes to keep social bots and fake content off their platforms. Twitter intensified its approach to this in mid-2018; the *Washington Post* reported that more than 70 million accounts were deleted for suspicious behavior in May and June 2018

alone (cf. Timberg/Dwoskin 2018). Facebook also published figures on deleted accounts for the first time in a transparency report in 2018, in which it stated that it had deleted more than a billion accounts in just six months: »694 million in the last quarter of 2017 and 583 million in the first quarter of 2018« (Brühl 2018).

Identifying social bots on social networks is a huge challenge for researchers. Social bots benefit enormously from behaving inconspicuously and imitating human behavior as accurately as possible in order to expand their influence. Often, the majority of available resources are invested in concealing the bot's identity. One of the core problems is the huge diversity of social bots used, meaning that different approaches for different types of social bot currently enjoy different levels of success. Another difficult factor is the fact that different data is available for investigation depending on the platform and the user's privacy settings: »[...] the majority of Twitter profiles are public, whilst on Facebook, most profiles are private« (Haugen 2017: 27). Furthermore, social bots are becoming ever more complex, putting them in a kind of arms race with researchers.

Before the content analysis itself began, this investigation used two freely available automated processes for bot detection: *Botometer* and *DeBot*. These use different strategies:

Botometer (originally published under the name BotOrNot) is a framework for detecting bots on Twitter that emerged from a collaboration between the Indiana University Network Science Institute (IUNI) and the Center for Complex Networks and Systems Research (CNetS). The framework uses more than 1,000 features from Twitter metadata, content, and interaction patterns in its classifications (cf. Davis et al. 2016). The features can be classified in six categories (cf. Varol et al. 2017):

1. User-based features: Meta-information such as the number of accounts followed and followers, the number of tweets posted, profile descriptions, and settings.
2. Friend-based features: Various language-related (e.g. entropy of the language, number of languages used) and temporal (e.g. age of the account) aspects, as well as the popularity of the tweet and the time between tweets, are considered based on the four types of information exchange (retweeting, mentioning, being retweeted, being mentioned).
3. Network-based features such as the retweet network, user mention network, and hashtag (coexistence) network.
4. Temporal features: The average rates of tweet production, retweets, and user mentions are considered based on measurements of activity at certain intervals.
5. Content and language-based features using statistical examinations of the length and information density of tweets.
6. Sentiment-based features: Information on the emotions communicated by a text.

In 2016, a team of three researchers from the University of New Mexico developed a totally different approach to bot detection – *DeBot*. As it does not analyze any explicit metadata from users, it is ideal for detecting bots that are part of a bot network. Instead, activity logs are looked at and examined for correlation. The researchers' fundamental theory is that »humans cannot be highly synchronous for a long duration; thus, highly synchronous user accounts are most likely bots« (Chavoshi et al. 2016). What makes this method so special is its ability to record even activity logs that have a time delay but are still synchronous.

Investigation design

In order to deduce possible patterns for bot use, the investigation combined technical detection processes and a content analysis, also based on algorithms, of the content produced by social bots and humans. The following research questions formed the basis of the investigation:

1. How high is the rate of social bots and what total proportion of the data sets investigated does their content account for?
2. Does the focus of the data set's content (politics, social affairs, consumption, lifestyle) have a significant influence on the rate of social bots?
3. How wide is the reach of the bot accounts found and what influence do the tweets spread by social bots have?
4. Does the content spread by social bots differ from human content in terms of text length, text mood, text subjectivity, media shared (images, videos, links) and the linking strategies used (hashtags, cashtags, user mentions)?

The data investigated consisted of tweets that were keyworded using a selection of hashtags over a defined investigation period of ten days. Each hashtag represents an exchange of information in various topic areas. All the tweets extracted in association with one hashtag were analyzed as independent data sets and the results compared at the end of the investigation. As Twitter itself deletes bots it identifies, it was only possible to investigate published tweets.

The data base

The period August 1, 2018 to August 10, 2018 was defined as the investigation period. This period of ten days was based on a rough estimate of the expected number of tweets given the research parameters chosen.

The hashtags and the associated topic areas of the content investigated were selected based on the following criteria:

- The number of tweets in the selected period must be sufficiently high (at least 10,000), but must not exceed 250,000 tweets in total (due to existing Twitter API limitations).
- The tweets under investigation should represent different content areas (politics, social affairs, consumption, lifestyle) in order to cover as many potential interests for the use of social bots as possible, e.g. influencing political opinion making, defaming political opponents, influencing discourse in society, shifting commercial interests etc.
- Due to the planned sentiment analysis, the tweets need to be in a single language. Based on the higher number of tweets in individual debates and the better compatibility with relevant frameworks, this language will be English. This means that only international hashtags and/or national hashtags in English-speaking countries should be selected.

Following detailed research into hashtags, the following four emerged as suitable:

#midTerms2018: The hashtag relates to the midterm elections in the USA held on November 6, 2018, in which the entire House of Representatives and a third of the Senate were up for election. The elections are seen as a test of the national mood ahead of the upcoming presidential election. The hashtag includes political discussions, links to media reporting, and opinion polls. Because the elections were in the near future at the time of the investigation, the hashtag also saw a lot of traffic during the investigation period.

#metoo: In October 2017, multiple allegations against the producer Harvey Weinstein became public, accusing him of the sexual abuse and sexual harassment of multiple women in the film industry. Shortly afterwards, Tarana Burke launched a campaign aiming to bring sexual assault and harassment into the open (cf. Göbel/Bäuerlein 2017). This campaign and the debate that accompanies it have since taken place under the hashtag #metoo, which also acts as a synonym for the MeToo movement of the same name. The hashtag has lost none of its topicality since its establishment, making it the ideal starting point for the planned analysis.

#iphone: In this selection, the hashtag #iphone predominantly represents content of a commercial or consumption-oriented nature and is thus intended to highlight a different focus than the other hashtags chosen. With the launch of the new iPhone model (iPhone XS) at around the same time, September 14, 2018, the hashtag saw particularly frequent use during the period selected due to various speculation and promotion.

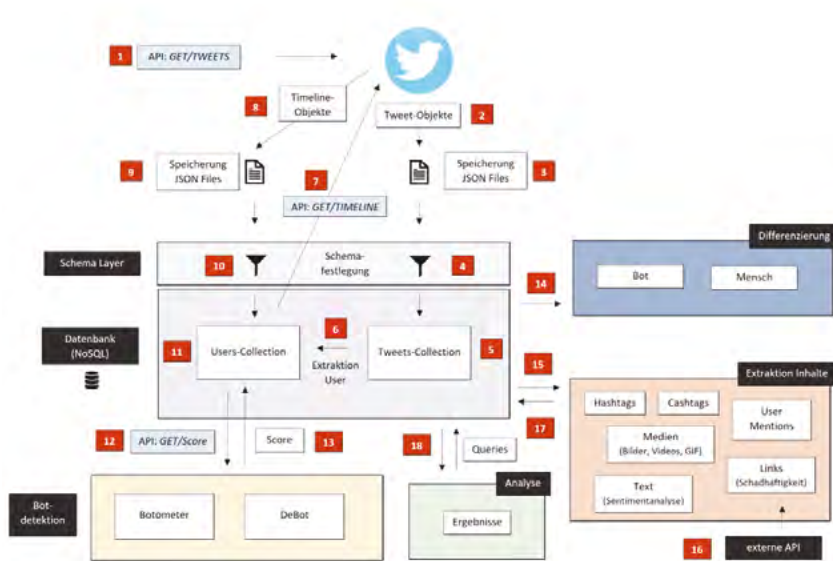
#foodporn: The hashtag #foodporn is intended to pick up on a very separate trend from the last few years – the (aesthetic) depiction of food on social media. Although this hashtag is related to the field of consumption, it is much more

focused on the lifestyle sector and adds a new topic area to the other three hash-tags.

Technical procedure

A diagram of the technical procedure for the investigation is shown in the figure below. Steps 1-16 have to be repeated separately for each data set under investigation (#midTerms2018, #metoo, #iphone, and #foodporn) in order to achieve comparable results at the end.

Figure 1
Steps of the investigation



Results: fewer bots, less influence than expected

Social bots performed worse than humans in every area: They had around 33 percent fewer followers on average, 63 percent fewer likes, 49 percent fewer retweets, and 67 percent fewer replies. Content comparisons were also conducted regarding the linking strategies used, incorporation of media, and external links and text

characteristics (sentiment analysis). The results of this do show some differences – social bots use around 50 percent more hashtags, 22 percent more links, and 20 percent more media, and produce texts that are around 30 percent shorter with slightly more negative and objective characteristics. However, no targeted influencing strategy was found.

The majority of research up to now has concentrated exclusively on political debates. The results of this paper suggest, however, that it would be useful to investigate content with a commercial background as well, in order to gain a more comprehensive view of the issue of social bots.

(1) How high is the rate of social bots and what total proportion of the data sets investigated does their content account for?

In total, 9,644 of 125,610 accounts investigated were identified as social bots – a rate of 7.68 percent. However, this rate differed widely between the various data sets and focus topics.

(2) Does the focus of the data set's content (politics, social affairs, consumption, lifestyle) have a significant influence on the rate of social bots?

Somewhat surprisingly, the lowest rate of social bots – 4.53 percent – was found in the political debate under #midTerms2018. Representing a debate in society, the #metoo discussion had a similarly low social bot rate of 6.60 percent. In contrast, lifestyle and consumption-related content appears to be much more attractive for the use of social bots, with 10.91 percent of the data sets under #foodporn and 17.79 percent of those under #iphone identified as social bot accounts.

(3) How wide is the reach of the bot accounts found and what influence do the tweets spread by social bots have?

At the same time, the investigation showed that social bots were comparatively productive. 13.18 percent of the 207,687 tweets investigated in total came from social bots. That puts them at almost twice as many tweets per account as humans, albeit with lower potential for influence. Social bots had around 33 percent fewer followers on average. In addition, content from humans had around two-and-a-half times more likes, twice as many retweets and three times as many replies compared to content from social bots. The #midTerms2018 data set was the only exception, with social bots achieving a range that was around 75 percent wider.

(4) Does the content spread by social bots differ from human content in terms of text length, text mood, text subjectivity, media shared (images, videos, links) and the linking strategies used (hashtags, cashtags, user mentions)?

Social bots tend to use hashtags as their preferred linking mechanism, employing them around 50 percent more frequently and with greater variance. Humans,

on the other hand, prefer user mentions as a way to address users directly, employing them more than twice as often.

Another finding relates to the use of external links. As found in studies by Stieglitz et al. 2017, Gilani et al. 2017, Brachten et al. 2017 and Chu et al. 2010, social bots use external links more frequently (by around 22 percent). Although the proportion of defective links was slightly higher for social bots, it remained low: 1.86 percent for social bots and 0.46 percent for humans.

Similar to the rate of links, media use was also around 20 percent higher for social bots – 0.55 media per tweet on average, compared to just 0.45 media per tweet for humans. When it comes to the distribution of the types of media (image, video, SIF), no significant difference was found between social bots and humans: images were by far the most popular medium among both groups.

In contrast to the greater use of links and media, the texts for social bots were around 30 percent shorter than those written by humans. The sentiment analysis of the texts did not provide any fundamentally different results, with the basic mood for both slightly positive. In a value range of -1/+1, the value for humans was +0.21 (μ : 0.29) and for social bots +0.14 (μ : 0.28). The data set for #midTerms2018 was an exception to this, with social bots delivering more negative texts compared to texts by humans. This could indicate a strategy of negative influence.

Although comparing the content did throw up some differences between social bots and humans, these differences were not significant and do not point to any explicit strategy. The inherent tactic of social bots – to communicate in as authentically human a way as possible – appears to have worked, at least formally. The lower influence of individual tweets combined with the higher tweet production suggests that social bots focus more on quantity than quality. The lower reach (33 percent fewer followers) can be seen to indicate lower influence.

The results of this investigation thus contrast significantly with reports in the media. Current research shows that social bots do not have excessive influence and are not disproportionately represented in social networks. Given the lack of research up to now, the true effect of the content spread by social bots on opinion forming and users' offline behavior also appears questionable.

What the media could learn from this

It would be easy to accuse the media of agenda setting. The causes can only be the subject of speculation. Do they lie in the insecurity of media creators or the ignorance of how algorithms work? Are they journalistic processes of news selection – when some warn against it, others pick up on it? Are there economic reasons behind the exaggeration of the potential dangers of social bots?

It is undoubtedly possible to conceive of theoretical dangers and risks resulting from social bots – from manipulative influencing of purchase choices and social discourse to influencing voting choices in elections. However, current findings do not provide clear evidence of their effectiveness. Some authors even speak of a »myth of social bots« (Gallwitz/Krell 2019).

Instead, one-sided reporting itself risks creating a distorted perception of social bots, which could lead to restrictive political consequences such as stricter regulation of social networks. This is the view taken by Linus Neumann, Spokesperson of the Chaos Computer Club. He sees social bots as a symptom rather than the cause of current developments in society (cf. Kind et al. 2017: 57). According to him, very different actors actually hold the potential for influence: »private television, the Bild newspaper, and lying Interior Ministers have a much greater political influence on people« (Rebiger 2017).

One thing is certain when it comes to media reporting: It should »promote conscious handling of the channels and discourses« – after all, »even algorithms make mistakes« (Niekler 2019). While some commentators warn of the power of social bots and call for Europe-wide regulation (Sarovic 2019), others see a tightening of the laws as a danger in itself. Markus Reuter shares the view that over-regulation presents a much more significant risk to democracy than the threat from bots, arguing: »The problem with fighting social bots or other manipulative accounts is that regulation very quickly begins to affect basic rights like freedom of speech and of the press – the negative consequences of regulation thus weigh heavier than the (negligible) damage to democracies currently identified from these disinformation tactics« (Reuter 2019).

This paper is based on a master's thesis in Media Informatics at Leipzig University of Applied Sciences by Tommy Hasert.

About the authors:

Tommy Hasert (*1988) gained his Master of Science in Media Informatics at Leipzig University of Applied Sciences. He works as a programmer in Leipzig. Contact: t.hasert@posteo.de

Gabriele Hooffacker (*1959), Dr.-phil., is a Professor teaching at the Faculty of Computer Science and Media at Leipzig University of Applied Sciences. She is co-editor of *Journalistik*. Contact: g.hooffacker@link-m.de

Translation: Sophie Costella

Literatur

- Abokhodair, Norah; Yoo, Daisy; McDonald, David W. (2015): Dissecting a Social Botnet. In: Dan Cosley, Andrea Forte, Luigina Ciolfi und David McDonald (Eds.): *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing - CSCW '15, the 18th ACM Conference*. Vancouver, BC, Canada, 14.03.2015 - 18.03.2015. New York, USA: ACM Press, S. 839–851
- Amann, Melanie; Knaup, Horand; Müller, Ann-Katrin; Rosenbach, Marcel; Wiedmann-Schmidt, Wolf (2016): Social Bots. Wie digitale Dreckschleudern Meinung machen. In: *Spiegel Online*. Online verfügbar unter <http://www.spiegel.de/spiegel/social-bots-propaganda-roboter-verzerrenpolitische-diskussion-im-netz-a-1117955.html#ref=rss>, zuletzt aktualisiert am 24.10.2016 [10.12.2018]
- Brachten, Florian, Stieglitz, Stefan, Hofeditz, Lennart, Kloppenborg, Katharina, & Reimann, Annette (2017): *Strategies and Influence of Social Bots in a 2017 German state election - A case study on Twitter*. ArXiv, abs/1710.07562.
- Brühl, Jannis (2018): Das große Scannen. Facebook. In: *Süddeutsche Zeitung Online*. Online verfügbar unter <https://www.sueddeutsche.de/wirtschaft/Facebook-das-grosse-scannen-1.3982384>, zuletzt aktualisiert am 16.05.2018 [13.11.2018]
- Davis, Clayton, Onur Varol, Emilio Ferrara, Alessandro Flammini, and Filippo Menczer (2016). BotOrNot: A System to Evaluate Social Bots. In: *Proceedings of the 25th International Conference on World Wide Web Companion*, pp. 273-274. <https://doi.org/10.1145/2872518.2889302>
- Gallwitz, Florian; Krell, Michael: Die Mär von »Social Bots«. In: *Tagesspiegel*, 03.06.2019, zuletzt aktualisiert 05.06.2019, online verfügbar unter <https://background.tagesspiegel.de/die-maer-von-social-bots> [07.06.2019]
- Hasert, Tommy (2018): *Einfluss von Social Bots in sozialen Netzwerken am Beispiel von Twitter*. Unveröffentlichte Masterarbeit, Studiengang Medieninformatik, HTWK Leipzig.
- Haugen, Geir Marius Sætvenes (2017): *Manipulation and Deception with Social Bots: Strategies and Indicators for Minimizing Impact*. Master Thesis. Norwegian University of Science and Technology
- Heglich, Simon; Janetzko, Dietmar (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
- Irmisch, Anna (2011): *Astroturf. Eine neue Lobbyingstrategie in Deutschland?* Wiesbaden: VS Verlag für Sozialwissenschaften. Online verfügbar unter <http://dx.doi.org/10.1007/978-3-531-92890-6>
- Kerl, Christian (2017): »Social Bots« – eine Gefahr für die Demokratie. In: *Berliner Morgenpost*, 21.1.2017. Online verfügbar unter <https://www.morgenpost.de/>

- [politik/article209348349/Social-Bots-eine-Gefahr-fuer-die-Demokratie.html](#)
[1.6.2019]
- Kind, Sonja; Bovenschulte, Marc; Ehrenberg-Silies, Simone; Jetzke Tobias; Weide Sebastian (2017): *Social Bots. TA-Vorstudie*. Hrsg. v. VDI/VDE Innovation + Technik GmbH. Büro für Technikfolgen-Abschätzung beim Deutschen Bundestag
- Meiselwitz, Gabriele (Ed.) (2017): *Social Computing and Social Media. Human Behavior*. 9th International Conference, SCSM 2017, Held as Part of HCI International 2017, Vancouver, BC, Canada, July 9-14, 2017, Proceedings, Part I
- Misener, Dan (2011): Rise of the socialbots: They could be influencing you online. In: *CBC News*. Online verfügbar unter <https://www.cbc.ca/news/technology/rise-of-the-socialbots-they-could-beinfluencing-you-online-1.981796>, zuletzt aktualisiert am 30.03.2011 [05.11.2018]
- Niekler, Andreas (2019): »Auch Algorithmen machen Fehler«. Online verfügbar unter <https://www.mediaqualitywatch.de/nachrichten/quot-auch-algorithmen-machen-fehler-quot> [07.06.2019]
- Ratkiewicz, J.; Conover, M. D.; Meiss, M.; Goncalves, B.; Flammini, A.; Menczer, F. (2011): Detecting and Tracking Political Abuse in Social Media. In: *Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media*, pp. 279-304
- Rebiger, Simon (2017): Fachgespräch im Bundestag: Experten halten Einfluss von Social Bots für überschätzt. In: *netzpolitik.org*. Online verfügbar unter <https://netzpolitik.org/2017/fachgesprach-im-bundestag-experten-halteneinfluss-von-social-bots-fuer-ueberschaetzt/>, zuletzt aktualisiert am 02.02.2017 [12.12.2018]
- Reuter, Markus (2019): Social Bots: Was nicht erkannt werden kann, sollte nicht reguliert werden. In: *netzpolitik.org*. Online verfügbar unter <https://netzpolitik.org/2019/social-bots-was-nicht-erkannt-werden-kann-sollte-nicht-reguliert-werden/>, zuletzt aktualisiert am 9.5.2019 [1.6.2019]
- Reuter, Markus (2017): Regulierungsdauerfeuer gegen Fake News und Social Bots ohne empirische Grundlage. In: *netzpolitik.org*. Online verfügbar unter <https://netzpolitik.org/2017/regulierungsdauerfeuer-gegen-fake-news-und-socialbots-ohne-empirische-grundlage/>, zuletzt aktualisiert am 24.01.2017 [10.12.2018]
- Sarovic, Alexander (2019): »Eine echte Bedrohung für die Demokratie«. In: *Spiegel online*. Online verfügbar unter <https://www.spiegel.de/politik/ausland/europawahl-2019-anne-applebaum-ueber-fake-news-und-desinformationskampagnen-a-1267521.html>, 23.5.2019 [1.6.2019]
- Timberg, Craig; Dvoskin, Elizabeth (2018): Twitter is sweeping out fake accounts like never before, putting user growth at risk. In: *Washington Post*. Online verfügbar unter https://www.washingtonpost.com/technology/2018/07/06/Twitter-is-sweepingout-fake-accounts-like-never-before-putting-user-growth-risk/?utm_term=.a5fo8e55c2ea, zuletzt aktualisiert am 06.07.2018 [13.11.2018]

- Varol, Onur; Ferrara, Emilio; Davis, Clayton A.; Menczer, Filippo; Flammini, Alessandro (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
- Watts, Clint (2018): Artificial intelligence is transforming social media. Can American democracy survive? In: *The Washington Post*. Online verfügbar unter https://www.washingtonpost.com/news/democracypost/wp/2018/09/05/artificial-intelligence-is-transforming-social-media-can-american-democracy-survive/?utm_term=.dfb1169f9c8e, zuletzt aktualisiert am 05.09.2018 [10.12.2018]
- Weck, Andreas (2016): Wie Social-Media-Trends durch Bots manipuliert werden. In: *t3n*. Online verfügbar unter <https://t3n.de/news/social-media-trends-bots-694529/> [06.11.2018]
- Woolley, Samuel; Gorbis, Marina (2017): Social media bots threaten democracy. But we are not helpless. In: *The Guardian*. Online verfügbar unter <https://www.theguardian.com/commentisfree/2017/oct/16/bots-social-mediathreaten-democracy-technology>, zuletzt aktualisiert am 16.10.2017 [10.12.2018]
- Woolley, Sam; Howard, Phil (2014): *Bad News Bots: How Civil Society Can Combat Automated Online Propaganda*. Online verfügbar unter <http://techpresident.com/news/25374/bad-news-bots-how-civil-society-cancombat-automated-online-propaganda> [05.11.2018]