Disagree? You Must be a Bot! How Beliefs Shape Twitter Profile Perceptions

MAGDALENA WISCHNEWSKI, University of Duisburg-Essen, Germany

REBECCA BERNEMANN, University of Duisburg-Essen, Germany

THAO NGO, University of Duisburg-Essen, Germany

NICOLE KRÄMER, University of Duisburg-Essen, Germany

In this paper, we investigate the human ability to distinguish political social bots from humans on Twitter. Following motivated reasoning theory from social and cognitive psychology, our central hypothesis is that especially those accounts which are opinion-incongruent are perceived as social bot accounts when the account is ambiguous about its nature. We also hypothesize that credibility ratings mediate this relationship. We asked $N = 151$ participants to evaluate 24 Twitter accounts and decide whether the accounts were humans or social bots. Findings support our motivated reasoning hypothesis for a sub-group of Twitter users (those who are more familiar with Twitter): Accounts that are opinion-incongruent are evaluated as relatively more bot-like than accounts that are opinion-congruent. Moreover, it does not matter whether the account is clearly social bot or human or ambiguous about its nature. This was mediated by perceived credibility in the sense that congruent profiles were evaluated to be more credible resulting in lower perceptions as bots.

CCS Concepts: • **Human-centered computing** → **Empirical studies in HCI**.

Additional Key Words and Phrases: motivated reasoning, social bots, Twitter, credibility, partisanship, bias

ACM Reference Format:

Magdalena Wischnewski, Rebecca Bernemann, Thao Ngo, and Nicole Krämer. 2021. Disagree? You Must be a Bot! How Beliefs Shape Twitter Profile Perceptions. In *CHI Conference on Human Factors in Computing Systems (CHI '21)*, May 8–13, 2021, Yokohama, Japan. ACM, New York, NY, USA, 17 pages. <https://doi.org/10.1145/3411764.3445109>

1 INTRODUCTION

Today's social media landscape not only enables individuals to access a rich well of information and to connect with others but also allows for user-generated content to be published online. However, not all activity on social media platforms originates from human users. In recent years, concerns about the influence of automated accounts, so-called social bots, has risen. Social bots on social media are generally defined as automatic or semi-automatic accounts run by computer algorithms [18]. As they often mimic human behavior (e.g. posting content, "liking", and "retweeting"), their automated nature can go unnoticed by human users. Social bots have been associated with a plethora of different functions like copy-paste bots which post the same content multiple times to gain attention [44], amplifier accounts that boost particular sources by aggregating and repeating content [22], or fake followers to boost follower counts and popularity [11].

The influence of social bots has become a public but also an academic concern. Bots have been, for example, accused of steering discussions online to promote specific ideas, spread misinformation [53], engage in political astroturfing

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Association for Computing Machinery.

Manuscript submitted to ACM

[29], affecting the stability of financial markets [6] and endorse conspiracy theories [47]. Moreover, influential bot activity has been found in Japan's general election in 2014 [44], the Brexit referendum in 2016 [3], the U.S. presidential election in 2016 [4], and the French general election in 2017 [17]. Besides their role in election campaigns, it was found that social bot activity is linked to detrimental effects on public health like the promotion of e-cigarettes [1] and the promotion of anti-vaccination content [7].

Due to these detrimental effects, many efforts have been made to detect social bots online. Most solutions employ machine learning models for automated detection, which rely on account features such as posting behavior, timing, or social networking structures. For a review of detection techniques, see [28]. Another methodological approach suggests identifying social bots through coordinated behavior or co-retweeting activity [29]. An example of such a strategy is to identify accounts that share identical content within a short period of time. These automated detection solutions received comparatively much academic attention as they are the only way to understand the influence of social bots on a larger scale. However, they cannot explain how users understand and detect social bots. In this contribution, we want to, therefore, focus on the lesser attended users' side of social bot detection.

We deem this important due to two reasons. First, although recent findings show that 58% of news consumers (from a global sample) are concerned about fake content as well as fake actors, only 26% report to have actually come across fake content/actors [38]. This shows that users are aware of the existence of social bots but may not be confident to identify them, resulting in a perception of threat and uncertainty. Second, with an increased uncertainty about what is real and what is fake, users' overall trustworthiness in (political) online communication may decrease. To support users in the best possible way in detecting social bots and to increase users' overall self-efficacy, we first need to understand under which circumstances users perceive an account as a (political) social bot. However, research on this is scarce. Hence, in this contribution, we want to pose the following research question:

RQ: Which underlying psychological mechanisms influence social bot detection of users on Twitter?

Relying on findings from social and cognitive psychology as well as political psychology, we know that, when it comes to political communication, individuals' perception can be biased [13, 48]. The theory of motivated reasoning suggests that information, which is in line with prior-opinions and attitudes, is favored over information that contradicts prior-opinions and attitudes. This results in overconfidence as well as over-reliance when information is opinion-congruent but also in rejection and overcritical assessment of opinion-incongruent information [31, 48]. In this paper, we therefore apply motivated reasoning theory to understand the human perception of Twitter accounts. Concretely, we suggest that due to motivated reasoning users may "blindly" approve of accounts only because the account represents one's own opinion, while accounts that disagree with one's own opinion are rejected or discredited as social bots.

As a first focus of this paper, we want to briefly summarize prior results on users' knowledge about social bots, users' engagement with social bots, users' detection abilities, and users' acceptance of social bots. We then introduce motivated reasoning theory and connect it to social bot perceptions. In doing so, we hypothesize that users' bot detection ability is biased in a way that opinion-congruent bots are less likely and opinion-incongruent bots more likely to be detected whereas opinion-congruent human profiles are perceived as such but opinion-incongruent human profiles become less human. At the core of this paper, we present a pre-registered online experiment, investigating our hypothesis. Because in the past, social bots have been found to influence political events such as elections, we chose the context of political communication for our experiment. Similarly, we selected Twitter as a social medium as a substantial presence of social bots has previously been found on Twitter [10] and it is a popular social medium for political communication [25].

Our overarching objective is to contribute to a better understanding of underlying mechanisms that drive (political) social bot perception on Twitter. Once we understand these mechanisms better, users' perception and detection of bots can be assisted through media literacy interventions as well as assisting tools. While media literacy interventions could, for example, correct erroneous user perceptions, assisting tools can support users in the identification process itself. Considering the detrimental effects of social bots, it is important that users can correctly identify social bot accounts.

This paper makes the following contributions:

- An empirical investigation of human interaction with social bots which is grounded in social psychological theory
- A deeper understanding of the conditions under which users perceive an account as a social bot
- Guidance for developers of assisting tools which want to support users in the detection of social bots

2 THEORETICAL BACKGROUND

2.1 What users know about and how they engage with social bots

The stated purpose of some social bots is to engage with users. Results for how users interact with social bots are, however, mixed. Building on the Computers are Social Actors (CASA) paradigm [43], Edwards and colleagues [15], for example, wanted to know if social bots are perceived differently from human agents when communicating on Twitter. As CASA suggests, they found that individuals perceived social bots "as a credible source of information" (p. 374) and competent communicator. They also found that individuals showed no difference in the intention to interact with a human agent or a social bot.

These results contradict later findings by Murgia and colleagues [37] who investigated human-bot interaction on Stack Overflow, a question and answer site for developers. Their results indicate that, although communication of the human agent and social bot was identical, answers given by the social bot were less accepted and received more down-votes than answers by the human agent. This was especially prevalent when the bot made an erroneous statement, leading the authors to suggest that humans have a "low tolerance for mistakes by a bot" (p. 6). The two opposing results could be explained, however, by the varying functionality of social bots in the experiments. While Edwards and colleagues [15] used social bots to inform users, social bots in Murgia's et al. [37] study were used to engage with users by providing answers to user questions.

Both studies assume, however, that users know about social bots. Yet, are users on social media aware of social bots? Answers to that are scarce. One exception comes from the Pew Research Center [46]. The authors found that two-thirds of people in the USA had at least heard about social bots, although only 16% heard about them a lot, with younger individuals being more likely to have heard about them. Moreover, most people believed that social bots were used for malicious purposes (80%). The authors also asked those participants, who indicated to know about social bots, how confident they were to detect a social bot on social media. About one half of the participants were somewhat confident to confident that they would recognize a social bot account on social media, with younger individuals, again, being more confident than elder individuals. This indicates that, with 60% knowing of social bots and of those 50% feeling confident to detect social bots, only one third of people in the USA felt confident to detect social bots.

To support users in identifying, especially, malicious social bots, different countermeasures have been developed and tested concerning their usability. Through a user experience survey, Yang et al. [56] wanted to know how users engage with Botometer, one of the early machine learning bot detection tools available to the public as well as a commonly used tool for academic purposes [12]. If users find a Twitter account to be suspicious, they can enter the account name

into Botometer which will return a probability score. Their results showed that most users found Botometer scores easy to understand (80%). Moreover, the authors found that users were equally concerned about false positives (humans erroneously classified as bots) and false negative results (bots erroneously classified as humans) a concern that has echoed recently within academia as well (see [41]). However, the survey by Yang et al. [56] does neither answer why and how people selected accounts they wanted to test nor how they interpreted the results. In other words, their work leaves open how people detected a suspicious account and why they made the decision to verify it on Botometer.

Concerning the acceptance of social bots, Hong and Oh [21] investigated how self-efficacy in identifying social news bots and greater prevalence of social news bots increased users' acceptance. The authors found that with increased self-efficacy in identifying social news bots, social news bots become more acceptable. Likewise, the more prevalent news from social news bots were perceived, the more acceptable they became. As self-efficacy is a subjective concept, the results, however, do not tell whether individuals who indicated high self-efficacy in identifying could actually identify social news bots correctly.

Lastly, Jakesch et al. [24] investigated the trustworthiness of AI-mediated communication compared to human-to-human communication on AirBnB. They found that users trusted AI less than humans but only in mixed environments where both were present. Applying their findings to communication of social bots and humans on Twitter, we could expect that users trust social bots simply less than humans. However, we expect a different outcome. While this study does not investigate users' trust, we suggest that users evaluate social bots, indeed, similar to humans due to motivated reasoning. We elaborate on this in the next section where we connect social bot detection with motivated reasoning.

To conclude, while some first attempts have been made to understand how users interact with social bots, how many users know of social bots, how they apply assisting tools such as Botometer, and when users accept social bots, it remains unknown which underlying psychological mechanisms drive social bot detection. While some of the here presented research suggests that bots should be evaluated differently to humans, we propose in the next section that may not always be the case.

2.2 Motivated reasoning theory and credibility cues in human social bot detection

In this study, we provide an attempt to understand the processes behind human social bot detection. In doing so, we apply the theory of motivated reasoning to understand users' detection abilities and relate this to credibility cues. Motivated reasoning theory generally proposes that incoming information is processed in a way to maintain existing attitudes and beliefs and disregard, reject or downplay opposing views. For example, a person cheering for a specific soccer club will reject any criticism against this club but happily agree that one's own team is the best in the world. While this example portrays a rather harmless consequence of motivated reasoning, engaging in motivated reasoning can become harmful easily. Replacing the soccer club in the example with a conspiracy follower might result into persons harming themselves or others.

Different psychological mechanisms of motivated reasoning have been identified. Early investigations, for example, suggest that this bias is a result of different motivational states [9, 31]. Others have connected motivated reasoning to cognitive biases such as confirmation bias [39] or myside bias [35] as well as affect-driven, biased memory retrieval [49] and discrete emotions [55].

Although the causes for motivated reasoning are debated within social psychology, effects of it have been found in various disciplines such as science communication [45], political communication [13], information selection on social media [54], misinformation detection and processing [52] or fact-checking [14]. While motivated reasoning has been found in many domains, real-world repercussions of motivated reasoning have been mostly associated with political

decision making [42]. In one study, Bisgaard [5] found, for example, that while partisans are willing to except a decrease of economic conditions, responsibilities were, however, attributed in a way to protect one's loyalties. In the context of misinformation, results from Ecker and Ang [14] indicated that, even after being retracted, misinformation had continued influence when they were in line with previous attitudes.

In line with the existing literature on motivated reasoning, we suggest that social media users engage in motivated reasoning when perceiving social bots, resulting in an overconfidence and "blindness" towards social bots that promote favored content, and skepticism and rejection towards bots promoting opposing views. We assume, however, that this motivated social bot detection is limited to accounts which are not clearly distinguishable as human- or bot-run (ambiguous accounts). If accounts can clearly be identified as human or bot (unambiguous accounts), we suggest that users do not engage in motivated bot detection, even if an account endorses an attitude contrary to one's own because individuals "draw the desired conclusion only if they can muster up the evidence necessary to support it" [31]. In other words, we assume that, if an account shows undeniable signs of automation or humanness, users cannot reason themselves into believing it is not automated or automated, respectively. We hypothesize:

Hypothesis 1

- a For accounts that are ambiguous whether they display a human or a bot account (ambiguous profiles), opinion-congruent accounts are more likely to be identified as human accounts.
- b For accounts that are ambiguous whether they display a human or a bot account (ambiguous profiles), opinion-incongruent accounts are more likely to be identified as bot accounts.

Hypothesis 2

- a For accounts that are clearly human (unambiguous human profiles), there is no differences between congruent and incongruent profiles concerning the correct identification as human accounts.
- b For accounts that are clearly bot (unambiguous bot profiles), there is no differences between congruent and incongruent profiles concerning the correct identification as bot accounts.

In addition to this, we argue that perceived account credibility will influence users' perception. It was found that users on social media employ a number of cues, specific to the respective platform, to arrive at a credibility evaluation [36]. Cues can concern the *source*, such as perceived knowledge, passion, transparency, reliability and influence as well as the *message* or *content*, such as consistency, accuracy and authenticity [27]. In the semantics of social media platforms, these cues have been translated into features such as metadata, follower/followee ratio, linguistics, use of hashtags, or visual- and image-based features [26]. Another strand of research examines credibility cues related to *posting behavior*, meaning, for example, temporal patterns of tweeting and retweeting [8]. In other words, accounts were perceived as less credible when posts occurred regularly and did not follow a circadian rhythm. However, in recent years, bots have become more sophisticated in their efforts to disguise their automated nature [2].

In connection with motivated reasoning theory, we argue that the use of the described credibility cues is biased in favor of one's own attitudes and beliefs. Hence, perceived credibility depends on the displayed opinion of an account. More credible accounts should be perceived as more human, whereas less credible accounts should be perceived as more bot-like. We hypothesize:

Hypothesis 3

The relationship between opinion-congruency and account perception is mediated by the perceived account credibility.

All hypotheses were pre-registered under <https://osf.io/97mcr/>. Hypothesis 3 was erroneously preregistered as a moderation hypothesis which was corrected in the online registration.

3 METHOD

A report on how we determined our sample size, data exclusions, manipulations, pre-test results and measures as well as the dataset of this study are available online (<https://osf.io/36mkw/>). The study received ethical approval by the ethics committee of the University of Duisburg-Essen.

3.1 Sample

For the online experiment, we recruited participants through Prolific.co, a UK-based online survey platform similar to Amazon Mechanical Turk. The experiment was conducted in late July 2020. Participation was restricted to U.S. Americans. The final sample consisted of $N = 151$ participants (60 male, 84 female, 6 non-binary, 1 not shared) with an age range from 18 to 68 years ($M = 28.86$, $SD = 10.18$). A majority of participants held at least a High School degree ($n = 72$) or a Bachelor's Degree ($n = 54$). Most participants identified White ($n = 86$), followed by Asian American ($n = 28$) and Black or African American ($n = 16$).

3.2 Experimental design and procedure

To test our hypotheses, we conducted a within-subject design with two independent factors, *ambiguity* and *congruency*. The factor ambiguity consisted of three levels "human", "social bot" and "ambiguous", specifying whether a Twitter account is ambiguous about its nature (neither clearly social bot nor clearly human) or not. The factor congruency consisted of the two levels, "opinion-congruent" and "opinion-incongruent" (see more on how we derived the respective levels in section 3.4).

The general task of participants was to identify whether a Twitter account represented a human-run account or a social bot. Before entering this task, we provided participants with a working definition of social bots¹ to ensure that all participants held a similar understanding. After that, we randomly presented 24 Twitter profiles to participants and asked them if they thought the profiles represented a human or social bot as well as how credible they found the profiles.

3.3 Dependent variables and control variables

After viewing each Twitter profile, participants were asked to indicate on a continuous sliding scale from 0 to 100 in whole integers, if they thought the profile represented a bot or a human run account, with values lower than 50 indicating a bot account and values greater than 50 a human account. The stronger it was perceived as a social bot, the lower the value and the stronger it was perceived as a human account, the higher the value. This constituted the dependent variable *profile perception*.

Participants were then asked to indicate how credible they found the displayed profile. To assess credibility measures, we used the 'trustworthiness' subscale from McCroskey and Teven's [33] source credibility scale. Credibility was measured through six semantic descriptions with two antonyms on each side of a 7-point Likert scale (honest/dishonest;

¹Working definition that was shown to participants before entering the detection task: "Social bots are automated online accounts that communicate more or less autonomously. They typically operate on social media sites, such as Twitter. Social bots serve different functions. Some are programmed to forward, like or comment on specific topics like the weather, sports results, but also political issues. Some others are programmed by companies to disseminate advertisements within social networking sites. Some bots state openly that they are automated accounts. Other accounts do not disclose their automated nature and can hardly be differentiated from human run accounts."

untrustworthy/trustworthy (R); honorable/dishonorable; moral/immoral; unethical/ethical (R); phony/genuine (R)). The sub-scale achieved a satisfactory Cronbach's $\alpha = .94$.

Because other variables might also contribute to the perception of social bots, we included several control variables in our study. These consisted of standard demographic data like age, gender, and education. We also included a measure for hours spent on social media per day ("Less than 1 hour", "1-3 hours", "4-7 hours", "more than 8 hours"), number of actively used social media platforms, whether Twitter was actively used ("Yes", "No"), whether participants knew about social bots ("Yes", "No") as well as media literacy. To control for media literacy, we used the subscale critical consumption developed by Koc and Barut [30] which consists of eleven items that are answered on a 5-point Likert scale, ranging from 1 = strongly disagree to 5 = strongly agree. The sub-scale achieved a satisfactory Cronbach's $\alpha = .87$.

3.4 Stimulus material and independent variables

For the first factor, *ambiguity*, we built the displayed Twitter profiles in a way that they were either unambiguous or ambiguous about their nature (human- or bot-run). To strengthen external validity, we first searched within Twitter for accounts that self-identified as social bots or that scored higher than .8 in Botometer. For human-run accounts, we selected verified profiles as examples of how to build genuine human-run accounts.

While we could not assure that selected social bots profiles were indeed social bots, we additionally focused on profiles that corresponded with previously found characteristics determining credibility perceptions. We selected three main characteristics that have previously been found to alter individuals' credibility perception. These characteristics were timing of posting behavior, shared content, and the displayed profile picture [16, 19, 26, 34]. Concerning the timing of posting behavior, we selected profiles in a way that postings were either frequent (e.g. post appeared every two minutes) or infrequent. For the shared content, we selected profile that posted (a) tweets only, (b) retweets only, (c) shares links only, or (d) shares a mix of a, b, and c. Concerning the profile picture, we either showed a picture of a person or a picture displaying a graphic (non-human). All profiles were mock-ups and did not represent real accounts on Twitter. This resulted in overall 15 mock-up Twitter profiles which were pretested concerning their credibility. All mock-up profiles were created in a way that they shared non-political, neutral content (e.g. content on animals or nature) to avoid the hypothesized bias that accounts become more or less bot-like because of opinion-congruency.

Results of the pre-test, which can be found in the supplementary material, indicated that three profiles were identified as social bots, six were identified as human accounts, and seven as ambiguous profiles, neither clearly bot nor clearly human. In a next step, we created unambiguous and ambiguous profiles. For the unambiguous profiles, we selected the three identified bot accounts and three of the identified human accounts to build new Twitter profiles. For each of these overall six accounts we created two version: one representing a Democrat account and one representing a Republican account. To increase external validity, we selected content that was shared by Democrats and Republicans on Twitter. To generate ambiguous profiles, we selected six of the pretested profiles that were neither classified as clearly human nor bot and created again two version: one representing a Democrat account and one representing a Republican account. An overview of the final set of created Twitter accounts is given in Table 1. One example for an unambiguous (social bot) profile is given in Figure 1. Each profile consisted of ten posts. All created Twitter profiles are shared in the supplementary material and can be viewed under: <https://osf.io/36mkw/>.

For the second factor, congruency, we measured individual partisanship through the 4-item Partisan Identity Scale developed by Huddy et al. [23]. The factor was then calculated by adding participants' partisanship for congruent evaluations (congruent = Democratic Twitter profile for self-identified Democrats / Republican Twitter profiles for

Fig. 1. Example for a neutral bot profile from the pre-test (left) and the final democrat bot profile used in the study (right). Here, we manipulated the timing of posts, the content of posts (non-sensical) as well as the sharing of only original tweets (no retweets).

Table 1. Overview of all created Twitter profiles.

	Unambiguous		Ambiguous
	Bot	Human	
Democrat	3	3	6
Republican	3	3	6

self-identified Republicans) and incongruent evaluations (incongruent = Republican Twitter profiles for self-identified Democrats / Democrat Twitter profiles for self-identified Republicans).

3.5 Strategy of analysis

In order to answer our hypotheses, we conducted a 2 x 3 repeated measure ANOVA. We originally pre-registered a 2 x 2 repeated measure ANOVA with only two levels of ambiguity. The distinction into three, instead of two, levels enabled us, however, to differentiate between human and social bot accounts. A detailed reasoning for a changed strategy of analysis can be found here as well as the results of the originally preregistered analysis (<https://osf.io/36mkw/>).

For hypothesis 3 (hypothesizing that the relationship between congruency and account perception is mediated by perceived account credibility), we ran mediation analyses for all three levels of ambiguity, using the PROCESS macro Version 3 by Hayes [20] for SPSS. We entered a grouping variable representing the two levels of opinion-congruency as a predictor into the model, the perceived credibility ratings as a mediator and the Twitter profile perception as the outcome variable.

Table 2. Correlation coefficients of all control variables correlated with each other. NML = New Media Literacy, SM = Social Media. ** $p < .01$, * $p < .05$.

	1. Age	2. NML	3. SM Platforms used	4. Time spent on SM	5. Social bot knowledge	6. Twitter usage
1.	1	-.004	-.17*	-.22 * *	.01	-.03
2.		1	.14	.11	-.14	.23**
3.			1	.21**	-.15	.44**
4.				1	.07	.24**
5.					1	-.16*

Table 3. Means and standard deviations of the dependent variable profile perception for the three levels of ambiguity. Values closer to 100 indicate higher perceived "humanness", while values closer to 0 indicate higher perceived "botness".

Ambiguity	Congruency	<i>M</i>	<i>SD</i>
Bot Profile	Congruent	39.7	20.8
	Incongruent	31.6	20.4
Human Profile	Congruent	76.6	17
	Incongruent	70.4	19.8
Ambiguous Profile	Congruent	57.4	16.2
	Incongruent	52.4	17.4

4 RESULTS

All analyses were conducted using R Studio Version 3.5 as well as SPSS Version 26 and pre-registered (<https://osf.io/97mcr/>).

4.1 Descriptive results

In addition to the demographic data described in the method section, we assessed five additional control variables: time spent on social media per day, number of social media platforms actively used, Twitter usage, social bot knowledge and media literacy. We found that two thirds of the participants in our sample spend between >1 and up to three hours a day on social media, while, on average, actively using 3-4 different social media platforms. Roughly two thirds reported to use Twitter actively, about 86% of the participants knew about social bots, and participants reached a self-reported mean media literacy of 3.9 ($SD = 0.58$) on a five-point Likert scale. Correlations of the control variables can be found in Table 2.

However, we wanted to know which accounts would be identified as social bots and whether participants would identify accounts in a biased manner (H1), assuming opinion-congruent accounts to be perceived as more human. Descriptive results of the dependent variable profile perception by condition are reported in Table 3. A visual analysis through density distributions of the dependent variable profile perception is presented in Figure 2.

According to the visual analysis reported in Figure 2, we found the results of our pre-test confirmed. Our predefined Twitter profiles classified as social bot accounts were identified as such (means below 50), whereas accounts designed to represent humans were identified as human (means above 50). Similarly, accounts designed to be ambiguous were perceived as such (means close to 50).

Fig. 2. Density distribution of the dependent variable *profile perception*. If a Twitter profile was perceived as a social bot, values are below 50. If the Twitter profile was perceived as a human profile, values are above 50.

4.2 Hypothesis testing through repeated measure ANOVA

As described in the section *strategy of analysis*, we conducted a 2 x 3 repeated measure ANOVA with the factors congruency (two levels) and ambiguity (three levels) to answer hypothesis H1 and H2. Because the assumption of sphericity was not met for the factor ambiguity, we used the Greenhouse-Geisser sphericity correction for this factor.

Supporting the results of the visual analysis reported in Figure 2, we found a significant main effect for ambiguity, $F(1.76, 150) = 7.87, p < .001, \eta_p^2 = .05$ (using Greenhouse-Geisser correction). This meant that participants evaluated profiles of the three different levels (bot, human, and ambiguous), in fact, differently. This supported our operationalization of the three different levels. Additionally, we found a significant interaction of ambiguity with the control variable age ($F(1.79, 150) = 7.29, p = .001, \eta_p^2 = .05$). This meant that, depending on age, users could differentiate better or worse between different levels of ambiguity. After analyzing this result visually, we found that younger individuals performed worse in differentiating between the three levels.

Concerning the factor congruency, we found no significant main effect, $F(1, 150) = 0.72, p = .39, \eta_p^2 = .005$, which indicated that, when holding profile ambiguity constant, there was no significant difference between congruent or incongruent profiles, other than motivated reasoning would suggest. In turn, this implied that hypothesis H1a and H1b had to be rejected. However, when scrutinizing the results of the control variables, we noted that three variables

interacted significantly with the factor congruency: social bot knowledge ($F(1, 150) = 5.52, p = .02, \eta_p^2 = .04$), age ($F(1, 150) = 13.29, p < .001, \eta_p^2 = .09$), and time spent on social media ($F(1, 150) = 4.51, p = .03, \eta_p^2 = .03$). Through visual analysis, we assessed what this interaction meant for each level of ambiguity. For all levels (human, social bot, and ambiguous profiles), we found that individuals who knew about social bots showed a greater effect of motivated reasoning. In other words, people with prior knowledge of social bots rated opinion-congruent profiles as more human and opinion-incongruent profiles as more bot-like compared to people who did not know about social bots. Similarly, we found that the more time people spend on social media the greater the motivated reasoning effect. In other words, people who spend more time on social media rated opinion-congruent profiles as more human and opinion-incongruent profiles as more bot-like compared to people who spend less time on social media. For age, the results were less clear. Younger individuals showed more motivated reasoning only for human profiles. For bot profiles as well as ambiguous profiles age did not show a substantial effect.

Moreover, we did not find a significant interaction between congruency and ambiguity, $F(1.86, 150) = 1.46, p = .23, \eta_p^2 = .01$ (using Greenhouse-Geisser correction), including the before mentioned control variables. This meant that our overall assumption that opinion-congruency would influence profile perceptions only for ambiguous profiles was not met (the combination of H1 and H2).

4.3 Planned contrasts

Through paired sample t-test, we found a significant difference between the means of ambiguous, opinion-congruent ($M = 57.4$) and ambiguous, opinion-incongruent ($M = 52.4$) accounts, $t(150) = 3.48, p = .001, d = 0.28$. Hence, the result for ambiguous opinion-congruent accounts supported H1a which suggested that these accounts would be perceived as human accounts (i.e. $M > 50$)². H1b was, however, not supported. We hypothesized that ambiguous opinion-incongruent accounts would be perceived as social bot profiles (i.e. $M < 50$). This was not the case ($M = 52.4$). Concludingly, we found that opinion-congruent accounts were perceived as *relatively more human* than opinion-incongruent accounts, supporting the motivated reasoning hypothesis.

To test H2a and H2b, we performed a paired sample t-test, adjusting for multiple comparisons through Bonferroni correction. We found neither support for H2a nor H2b. Means of opinion-congruent ($M = 39.7$) and opinion-incongruent ($M = 31.6$) social bots accounts were significantly different, $t(150) = 4.72, p < .001, d = .38$. Similarly, means of opinion-congruent ($M = 76.6$) and opinion-incongruent ($M = 70.4$) human accounts were significantly different as well, $t(150) = 3.77, p < .001, d = .3$. In both cases, opinion-congruent accounts were perceived as relatively more human, similarly to ambiguous accounts. However, we hypothesized that accounts that are clearly identifiable as human or bot should be perceived as such, independent of the displayed opinion in the account. We conclude that the effect of motivated reasoning is stronger than we originally anticipated as we found motivated reasoning even for accounts that were previously identified as clearly human and social bot.

4.4 Mediation analysis for credibility

We conducted overall three mediation analysis, representing the three levels of ambiguity and following the procedure described in the strategy of analysis. The outcome variable for each analysis was the perception of the Twitter profiles,

²The results of the planned contrasts might seem to contradict findings of the previously reported repeated measure ANOVA. Yet, readers should bear in mind that the paired sample t-test did not include control variables. The above reported results of the repeated measure ANOVA give a more refined insight into the data.

whereas the two levels of congruency were the predictor variable. The mediator variable was the credibility rating of the respective level of ambiguity.

We tested the significance of this effect using bootstrapping procedures, computing 5000 bootstrapped samples with a confidence interval of 95%. The unstandardized indirect effect coefficient of credibility of congruent and incongruent social bot profiles was -7.59 with a 95% confidence interval ranging from -10.94 to -4.48 (see also Figure 3a). This supported our hypothesis that the perception of bot profiles is mediated by perceived credibility of the social bot profile. Similarly, we found a significant indirect effect of credibility of congruent and incongruent human profiles of -8.27 with a 95% confidence interval ranging from -11.12 to -5.76 (see Figure 3b) as well as a significant indirect effect of credibility of congruent and incongruent ambiguous profiles of -9.77 with a 95% confidence interval ranging from -12.64 to -7.11 (see Figure 3c). Overall, this results show that opinion-congruent profiles were perceived as more credible which lead users to rate them as more human, whereas profiles that were opinion-incongruent were perceived as less credible which lead users to rate them as more bot-like.

Fig. 3. Mediation analyses with relevant path coefficients (unstandardized) for (a) bot profiles, (b) human profiles, and (c) ambiguous profiles. Significant path coefficients are marked as followed: *** $p < .001$, ** $p < .01$, * $p < .05$.

5 DISCUSSION

In this study, we wanted to know under which conditions users perceive an account to be a social bot or a human. Following motivated reasoning theory, we hypothesized that Twitter profiles which are opinion-congruent would be more likely to be identified as human accounts (H1a), whereas opinion-incongruent accounts would be more likely to be identified as social bot accounts (H1b). We also suggested that this would only occur in cases where users are unable to differentiate between social bots and human accounts (ambiguous accounts). In other words, if profiles clearly represent a human (H2a) or a social bot (H2b) account, opinion-congruency has no effect on how the profile is perceived.

Additionally, we predicted that the relationship between opinion-congruency and profiles perception is mediated by perceived profile credibility.

Our results indicate that participants perceived the levels of Twitter profiles namely, human, social bot and ambiguous profiles, differently (a significant main effect for the factor ambiguity) as can be seen in Figure 3. The results of paired sample t-tests also supported this. Consequently, we observe that participants could generally differentiate between social bot and human accounts.

Moreover, we find that younger individuals differentiated less between the three levels of Twitter profiles as compared to older individuals, independent of the levels of opinion-congruency. We suggest that the increased exposure of younger individuals to social media (see significant correlations of age with time spent on social media and social media platforms used in Table 2) makes them more suspicious of accounts. This is why they are unable to ascertain the identity of any account (human or social bot). However, to understand why participants perceived a particular account as social bot or human can only be answered through in-depth qualitative studies.

5.1 The role of opinion-congruency when perceiving Twitter accounts

In relation to our central hypothesis, suggesting that profiles are perceived in a biased manner, we found only partial evidence to support this claim. When including control variables, we did not find a main effect for opinion-congruency. This indicates that opinion-congruent and opinion-incongruent profiles were perceived similarly. This result rejects our hypothesis that motivated reasoning guides profile perception. However, we found a motivated reasoning effect in combination with the control variables *age*, *time spent on social media* and *social bot knowledge*. We found that, especially participants with prior knowledge of social bots as well as participants who spend more time on social media, showed a bias. In line with motivated reasoning, they perceived opinion-congruent profiles as more human and opinion-incongruent profiles as more bot-like.

We propose that these results stem from a different usage of the term social bot. We argue that participants with prior knowledge of social bots and participants who spend more time on social media might apply the term social bot as a pejorative term to indicate disagreement and discredit accounts in comparison to less well-versed users. We believe that the effect of motivated reasoning is, in turn, amplified by wanting to show disagreement by labeling accounts as social bots (expressive disagreement). Popular media have discussed this labeling of others as social bots to discredit them before. In 2019, Sašo Ornik discussed in a Medium article how the political opposition in Slovenia, as well as the USA, have been demonized and degraded by associating them with social bots or an evil mastermind ³.

Moreover, when excluding the control variables from our analysis through paired-sampled t-tests, we found that mean profile perceptions of opinion-congruent profiles were, as predicted, rated relatively more human (H1a) whereas opinion-incongruent profiles were rated as relatively more bot-like (H1b). This was a constant pattern overall profiles, independent of the levels of ambiguity. This means that a motivated reasoning effect could only be found when not controlling for individual differences such as age, time spent on social media or knowledge of social bots.

Our results also reveal the expected mediation of congruency and profile perceptions through profile credibility. This means that perceived profile credibility contributes significantly to the effect of opinion-congruency on the profile perception whereby opinion-congruent profiles were perceived as more credible than opinion-incongruent profiles. We want to emphasize, however, that, similarly to the paired-sample t-tests, these results do not include control variables.

³<https://medium.com/@saosornik/everybody-i-dont-agree-with-is-a-russian-bot-or-how-it-is-easier-to-believe-an-evil-mastermind-ca02391055cb>

Concludingly, we want to come back to our research question, asking under which circumstances users perceive an account as a social bot. By applying motivated reasoning theory and credibility judgements, our results show that users' political opinion influences their ability to detect social bots but only for users with a high familiarity of social media.

5.2 Practical implications and future directions

Our findings pose challenges to those designing treatments against social bots as well as raise new research questions. Concerning designers and developers, we found empirical evidence that some user groups perceive profiles in a biased manner. For assisting tools like Botometer, which rely on users to self-identify accounts they want to check, this suggests that users might enter simply profiles they disagree with, expecting it to be a social bot, due to motivated reasoning. This implies that, while automated detection tools may improve their validity over time, users might apply these tools in a biased manner. Hence, empirical studies should examine whether accounts that are incongruent to one's own opinion are more likely to be tested of being a social bot. In turn, motivated reasoning also suggests that detection results from tools like Botometer might become less credible because they do not confirm one's expectation - a disconfirmation bias [48]. Future studies should investigate whether results of assisting tools are less credible to users, if they do not confirm the expected outcome. As there is no way for users to identify bots with certainty, we conclude that social media companies need to disclose the nature of an account. The responsibility of accurately identifying accounts or an unbiased use of detection tools should not rest on users.

Moreover, our findings reveal that those users, who are more familiar with social media, are more likely to show the hypothesized bias. Research on fluency - the ease with which an initial, intuitive answer is produced - suggests that familiarity increases the experience of fluency. In turn, high fluency increases *feelings of rightness* and retrospective confidence judgments [51]. The authors relate this to the use of deliberative processing. Accordingly, high fluency relates to less deliberative processing and, instead, evokes the reliance on heuristics and intuitive processing. In addition, previous studies have found that individuals who rely on heuristic processing show a greater motivated reasoning bias [40, 50]. Taken together, we hypothesize that individuals who are more familiar with social media environments experience greater fluency which results in less deliberative processing, leading to increased motivated reasoning. However, to gain a deeper understanding of these processes, future empirical studies should more systematically investigate the influence of familiarity with social media environments and bias. As we found a mediating effect for credibility, the effect of familiarity could also affect perceived credibility.

Furthermore, we suggest the inclusion of motivated reasoning and its effects in media literacy programs. This has been proposed before, for example, Lenker [32] argued that individuals need to be educated about the detrimental effects of motivated reasoning as well as to reflect upon one's own motivated reasoning when processing information. Future studies should investigate whether educating users about motivated reasoning results into less biased perception.

5.3 Limitations

Participants were only shown Twitter profiles. Limitations that come with this are threefold: First, our results are platform dependent. Although we asked whether participants actively used Twitter and controlled for this factor, we can neither make claims about specific other social media platforms, nor can we assess how generalizable our results are. Second, users were presented full Twitter profiles of only ten posts. Actual Twitter profiles go beyond ten posts and offer more information which can help better identify an account. Third, users are most likely to come across social bots on their Twitter timeline, where they can only view one post. However, one post is less indicative about an accounts authenticity than an account profile.

Lastly, the social bot accounts which we used in this study were artificially created. Although similar to social bots on Twitter as well as pre-tested, they represent only a limited type of social bots. The cues that we used to build the profiles do not represent all possible cues and combination of cues. For example, we did not manipulate follower/followee ratio which is a commonly used cue to assess whether an account is automated or not. Future studies would profit from assessing users' perception among a variety of different social media platforms, within different contexts (timeline versus profiles) and with different cues associated with social bots.

6 CONCLUSION

In this study, we investigated as to when users perceive an account to be a social bot or an human. In line with motivated reasoning, we hypothesized that opinion-congruent accounts are perceived as more human and opinion-incongruent accounts as more bot-like. Our main finding supports this hypothesis for a subgroup of users: users' ability to differentiate human and social bot accounts is affected by their prior-opinion. We conclude that participants who know about social bots and participants who spend more time on social media labeled profiles as social bots to express disagreement and discredit the profile. Moreover, we assert that opinion-congruent profiles were perceived as more credible than opinion-incongruent profiles.

ACKNOWLEDGMENTS

This research was funded by the German Research Foundation (DFG) under grant No. GRK 2167, Research Training Group "User-Centred Social Media".

REFERENCES

- [1] Jon-Patrick Allem, Emilio Ferrara, Sree Priyanka Uppu, Tess Boley Cruz, and Jennifer B. Unger. 2017. E-Cigarette surveillance with social media data: Social bots, emerging topics, and trends. *JMIR Public Health and Surveillance* 3, 4 (2017), e98. <https://doi.org/10.2196/publichealth.8641>
- [2] Eiman Alothali, Nazar Zaki, Elfadil A. Mohamed, and Hany Alashwal. 2019. Detecting social bots on Twitter: A literature review. *Proceedings of the 13th International Conference on Innovations in Information Technology, IIT 2018* (2019), 175–180. <https://doi.org/10.1109/INNOVATIONS.2018.8605995>
- [3] Marco T. Bastos and Dan Mercea. 2019. The Brexit botnet and user-generated hyperpartisan news. *Social Science Computer Review* 37, 1 (2019), 38–54. <https://doi.org/10.1177/0894439317734157>
- [4] Alessandro Bessi and Emilio Ferrara. 2016. Social bots distort the 2016 U.S. Presidential election online discussion. *First Monday* 21, 11 (2016), 1–17. <https://doi.org/10.5210/fm.v21i11.7090>
- [5] Martin Bisgaard. 2015. Bias will find a way: Economic perceptions, attributions of blame, and partisan-motivated reasoning during crisis. *The Journal of Politics* 77, 3 (2015), 849–860. <https://doi.org/10.1086/681591>
- [6] Johan Bollen, Huina Mao, and Xiaojun Zeng. 2011. Twitter mood predicts the stock market. *Journal of Computational Science* 2, 1 (2011), 1–8.
- [7] David A. Broniatowski, Amelia M. Jamison, Si Hua Qi, Lulwah AlKulaib, Tao Chen, Adrian Benton, Sandra C. Quinn, and Mark Dredze. 2018. Weaponized health communication: Twitter bots and Russian trolls amplify the vaccine debate. *American Journal of Public Health* 108, 10 (2018), 1378–1384. <https://doi.org/10.2105/AJPH.2018.304567>
- [8] Chiyu Cai, Linjing Li, and Daniel Zengi. 2017. Behavior enhanced deep bot detection in social media. In *2017 IEEE International Conference on Intelligence and Security Informatics (ISI)*. IEEE, 128–130.
- [9] Shelly Chaiken, Roger Giner-Sorolla, and Serena Chen. 1996. Beyond accuracy: Defense and impression motives in heuristic and systematic information processing. In *The Psychology of action: Linking cognition and motivation to behavior*, P. M. Gollwitzer and J. A. Bargh (Eds.). Guilford Press, New York, NY, 553–578.
- [10] Zi Chu, Steven Gianvecchio, Haining Wang, and Sushil Jajodia. 2010. Who is tweeting on Twitter: Human, bot, or cyborg?. In *Proceedings of the 26th annual computer security applications conference*. 21–30.
- [11] Stefano Cresci, Roberto Di Pietro, Marinella Petrocchi, Angelo Spognardi, and Maurizio Tesconi. 2015. Fame for sale: Efficient detection of fake Twitter followers. *Decision Support Systems* 80 (2015), 56–71. <https://doi.org/10.1016/j.dss.2015.09.003>
- [12] Clayton A. Davis, Onur Varol, Emilio Ferrara, Alessandro Flammini, and Filippo Menczer. 2016. BotOrNot: A system to evaluate social bots. (2016), 4–5. <https://doi.org/10.1145/2872518.2889302> arXiv:1602.00975
- [13] James N. Druckman, Thomas J. Leeper, and Rune Slothuus. 2016. Motivated responses to political communications: Framing, party cues, and science information. In *The Feeling, Thinking Citizens* (1 ed.), Howard Lavine and Charles S. Taber (Eds.). Routledge, New York, 125–150. <https://doi.org/10.1080/17445019.2016.1191111>

- [//doi.org/10.4324/9781351215947](https://doi.org/10.4324/9781351215947)
- [14] Ullrich K.H. Ecker and Li Chang Ang. 2019. Corrections, political attitudes and the processing of misinformation. *Political Psychology* 40, 2 (2019), 241–260. <https://doi.org/10.1111/pops.12494>
 - [15] Chad Edwards, Autumn Edwards, Patric R. Spence, and Ashleigh K. Shelton. 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 (2014), 372–376. <https://doi.org/10.1016/j.chb.2013.08.013>
 - [16] Richard M. Everett, Jason R.C. Nurse, and Arnau Erola. 2016. The anatomy of online deception: What makes automated text convincing?. In *Proceedings of the 31st Annual ACM symposium on applied computing*. 1115–1120. <https://doi.org/10.1145/2851613.2851813>
 - [17] Emilio Ferrara. 2017. Disinformation and social bot operations in the run up to the 2017 French presidential election. *First Monday* 22, 8 (2017). <https://doi.org/10.5210/fm.v22i8.8005>
 - [18] Emilio Ferrara, Onur Varol, Clayton Davis, Filippo Menczer, and Alessandro Flammini. 2016. The rise of social bots. *Commun. ACM* 59, 7 (2016), 96–104. <https://doi.org/10.1145/2872518.2889302>
 - [19] Christian Grimme, Mike Preuss, Lena Adam, and Heike Trautmann. 2017. Social bots: Human-like by means of human control? *Big Data* 5, 4 (2017), 279–293. <https://doi.org/10.1089/big.2017.0044> arXiv:1706.07624
 - [20] Andrew F Hayes. 2017. *Introduction to mediation, moderation, and conditional process analysis: A regression-based approach*. Guilford publications.
 - [21] Hyehyun Hong and Hyun Jee Oh. 2020. Utilizing bots for sustainable news business: Understanding users' perspectives of news bots in the age of social media. *Sustainability* 12, 16 (2020), 6515. <https://doi.org/10.3390/sui2166515>
 - [22] Philip N. Howard and Bence Kollanyi. 2016. Bots, #StrongerIn, and #Brexit: Computational propaganda during the UK-EU referendum: Research note 2016.1. *Oxford: Computational Propagand Project, 2016* (2016).
 - [23] Leonie Huddy, Lilliana Mason, and Lene Aarøe. 2015. Expressive partisanship: campaign involvement, political emotion, and partisan identity. *American Political Science Review* 109, 1 (2015), 1–17. <https://doi.org/10.1017/S0003055414000604>
 - [24] Maurice Jakesch, Megan French, Xiao Ma, Jeffrey T Hancock, and Mor Naaman. 2019. AI-mediated communication: How the perception that profile text was written by AI affects trustworthiness. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–13. <https://doi.org/10.1145/3290605.3300469>
 - [25] Andreas Jungherr. 2015. Twitter as political communication space: Publics, prominent users, and politicians. In *Analyzing Political Communication with Digital Trace Data. Contributions to Political Science*. Springer, Cham. https://doi.org/10.1007/978-3-319-20319-5_4
 - [26] Byungkyu Kang, Tobias Höllerer, and John O. Donovan. 2015. Believe it or not? Analyzing information credibility in microblogs. *Proceedings of the 2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining* 2015 (2015), 611–616.
 - [27] Minjeong Kang. 2010. Measuring social media credibility: A study on a measure of blog credibility. *Institute for Public Relations* (2010), 59–68. <https://doi.org/10.1136/bmj.g5133>
 - [28] Arzum Karataş and Serap Şahin. 2017. A review on social bot detection techniques and research directions. *Proc. Int. Security and Cryptology Conference Turkey* i (2017), 156–161.
 - [29] Franziska B. Keller, David Schoch, Sebastian Stier, and Junghwan Yang. 2020. Political astroturfing on Twitter: How to coordinate a disinformation campaign political astroturfing on Twitter. *Political Communication* 37, 2 (2020), 256–280. <https://doi.org/10.1080/10584609.2019.1661888>
 - [30] Mustafa Koc and Esra Barut. 2016. Development and validation of New Media Literacy Scale (NMLS) for university students. *Computers in Human Behavior* 63 (2016), 834–843. <https://doi.org/10.1016/j.chb.2016.06.035>
 - [31] Ziva Kunda. 1990. The case for motivated reasoning. *Psychological Bulletin* 108, 3 (1990), 480–498. <https://doi.org/10.1037/0033-2909.108.3.480>
 - [32] Mark Lenker. 2016. Motivated reasoning, political information, and information literacy education. *portal: Libraries and the Academy* 16, 3 (2016), 511–528. <https://doi.org/10.1353/pla.2016.0030>
 - [33] James C. McCroskey and Jason J. Teven. 1999. Goodwill: A reexamination of the construct and its measurement. *Communications Monographs* 66, 1 (1999), 90–103. <https://doi.org/10.1080/03637759909376464>
 - [34] Judith Meinert, Ahmet Aker, and Nicole C. Krämer. 2019. The impact of Twitter features on credibility ratings - An explorative examination combining psychological measurements and feature based selection methods. In *Proceedings of the 52nd Hawaii International Conference on System Sciences*. 2600–2609. <https://doi.org/10.24251/HICSS.2019.313>
 - [35] Hugo Mercier. 2016. Confirmation bias - myside bias. In *Cognitive Illusions: Intriguing Phenomena in Judgement, Thinking and Memory* (2 ed.), Rüdiger Pohl (Ed.). Psychology Press, London, 99–114. <https://doi.org/10.4324/9781315696935>
 - [36] Miriam J. Metzger, Andrew J. Flanagin, and Ryan B. Medders. 2010. Social and heuristic approaches to credibility evaluation online. *Journal of Communication* 60, 3 (2010), 413–439. <https://doi.org/10.1111/j.1460-2466.2010.01488.x>
 - [37] Alessandro Murgia, Daan Janssens, Serge Demeyer, and Bogdan Vasilescu. 2016. Among the machines: Human-bot interaction on social Q&A websites. *Conference on Human Factors in Computing Systems - Proceedings* 07-12-May- (2016), 1272–1279. <https://doi.org/10.1145/2851581.2892311>
 - [38] Nic Newman, Antonis Fletcher, Richard, Kalogeropoulos, and Rasmus Kleis Nielsen. 2019. Reuters Institute Digital News Report 2019. *Reuters Institute for the Study of Journalism* 2019 (2019), 1– 156. arXiv:arXiv:1011.1669v3
 - [39] Raymond S. Nickerson. 1998. Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology* 2, 2 (1998), 175–220.
 - [40] Gordon Pennycook and David G. Rand. 2018. Lazy, not biased: Susceptibility to partisan fake news is better explained by lack of reasoning than by motivated reasoning. *Cognition* (2018), 1–12. <https://doi.org/10.1016/j.cognition.2018.06.011>

- [41] Adrian Rauchfleisch and Jonas Kaiser. 2020. The false positive problem of automatic bot detection in social science research. *SSRN Electronic Journal* 7641 (2020). <https://doi.org/10.2139/ssrn.3565233>
- [42] David P. Redlawsk. 2002. Hot cognition or cool consideration? Testing the effects of motivated reasoning on political decision making author(s). *The Journal of Politics* 64, 4 (2002), 1021–1044. <https://doi.org/10.1111/1468-2508.00161>
- [43] Byron Reeves and Clifford I. Nass. 1996. *The media equation: How people treat computers, television, and new media like real people and places*. Cambridge University Press, New York.
- [44] Fabian Schäfer, Stefan Evert, and Philipp Heinrich. 2017. Japan's 2014 general election: Political bots, right-wing internet activism, and Prime Minister Shinz o Abe's hidden nationalist agenda. *Big Data* 5, 4 (2017), 294–309. <https://doi.org/10.1089/big.2017.0049>
- [45] Dietram A. Scheufele and Nicole M. Krause. 2019. Science audiences, misinformation, and fake news. *Proceedings of the National Academy of Sciences* 116, 16 (2019), 7662–7669. <https://doi.org/10.1073/pnas.1805871115>
- [46] Galen Stocking and Nami Sumida. 2018. Social media bots draw public's attention and concern. *Pew Research Center* October (2018).
- [47] Venkatramanan S. Subrahmanian, Amos Azaria, Skylar Durst, Vadim Kagan, Aram Galstyan, Kristina Lerman, Linhong Zhu, Emilio Ferrara, Alessandro Flammini, and Filippo Menczer. 2016. The DARPA Twitter bot challenge. *Computer* 49, 6 (2016), 38–46. <https://doi.org/10.1109/MC.2016.183>
- [48] Charles S. Taber and Milton Lodge. 2006. Motivated skepticism in the evaluation of political beliefs. *American Journal of Political Science* 50, 3 (2006), 755–769. <https://doi.org/10.1111/j.1540-5907.2006.00214.x> arXiv:9809069v1 [arXiv:gr-qc]
- [49] Charles S Taber and Milton Lodge. 2016. The illusion of choice in democratic politics: The unconscious impact of motivated political reasoning. *Political Psychology* 37 (2016), 61–85. <https://doi.org/10.1111/pops.12321>
- [50] Ben M. Tappin, Gordon Pennycook, and David G. Rand. 2020. Rethinking the link between cognitive sophistication and politically motivated reasoning. *Journal of Experimental Psychology: General* (2020). <https://doi.org/10.1037/xge0000974>
- [51] Valerie A. Thompson, Jamie A. Prowse Turner, Gordon Pennycook, Linden J. Ball, Hannah Brack, Yael Ophir, and Rakefet Ackerman. 2013. The role of answer fluency and perceptual fluency as metacognitive cues for initiating analytic thinking. *Cognition* 128, 2 (2013), 237–251. <https://doi.org/10.1016/j.cognition.2012.09.012>
- [52] Jay J. Van Bavel and Andrea Pereira. 2018. The partisan brain: An identity-based model of political belief. *Trends in Cognitive Sciences* 22, 3 (2018), 213–224. <https://doi.org/10.1016/j.tics.2018.01.004>
- [53] Patrick Wang, Rafael Angarita, and Ilaria Renna. 2018. Is this the era of misinformation yet: Combining social bots and fake news to deceive the masses. In *Companion Proceedings of the The Web Conference 2018*. 1557–1561. <https://doi.org/10.1145/3184558.3191610>
- [54] Stephan Winter, Miriam J. Metzger, and Andrew J. Flanagin. 2016. Selective use of news cues: A multiple-motive perspective on information selection in social media environments. *Journal of Communication* 66, 4 (2016), 669–693. <https://doi.org/10.1111/jcom.12241>
- [55] Nicole C. Wischniewski, Magdalena und Krämer. in press. *Technology, Mind and Behavior* (in press).
- [56] Kai Cheng Yang, Onur Varol, Clayton A. Davis, Emilio Ferrara, Alessandro Flammini, and Filippo Menczer. 2019. Arming the public with artificial intelligence to counter social bots. *Human Behavior and Emergent Technologies* (2019), 48–61. <https://doi.org/10.1002/hbe.2.115>