

“I Have No Idea What a Social Bot Is”: On Users’ Perceptions of Social Bots and Ability to Detect Them

ABSTRACT

Social bots—software-controlled accounts on online social networks (OSNs)—have been employed for various malicious purposes, including spreading disinformation and scams. Understanding users’ perceptions of bots and ability to distinguish them from other accounts can inform mitigations. To this end, we conducted an online study with 297 users of seven OSNs to explore their mental models of bots and evaluate their ability to classify bots and non-bots correctly. We found that while some participants were aware of bots’ primary characteristics, others provided abstract descriptions or confused bots with other phenomena. Participants also struggled to classify accounts correctly (e.g., misclassifying >50% of accounts) and were more likely to misclassify bots than non-bots. Furthermore, we observed that perceptions of bots had a significant effect on participants’ classification accuracy. For example, participants with abstract perceptions of bots were more likely to misclassify. Informed by our findings, we discuss directions for developing user-centered interventions against bots.

1 INTRODUCTION

Social bots have gained increased attention in recent years. These are typically automated or semi-automated accounts on online social networks (OSNs) controlled by computer software while potentially imitating humans. While bots are sometimes used for benign aims such as assisting users in reading long content [28], they are often deployed for malicious purposes, such as disseminating malware [19], participating in spam and phishing campaigns [48], spreading disinformation [32], popularizing conspiracy theories [21], and polarizing online conversations [51].

In return, researchers have invested significant efforts to understand the behavioral properties of bots, and the mechanisms that drive their flourishing ecosystem [1, 26, 35]. Guided by this understanding, researchers developed numerous bot-detection technologies to help moderate bots and counter their abuse [11, 17, 57]. These technologies, sometimes further supported by human experts, helped OSNs suspend many malicious bot campaigns [54].

Despite these efforts, malicious bots remain widely active on popular OSNs. Varol et al. estimated that 9%–15% of active Twitter accounts are bots [53]. More recently, a whistleblower reported that Facebook took no actions against bots used for political manipulation, despite being aware of their activity [9]. Thus, as the burden of detecting bots falls on OSNs’ users, there remains a need for improved interventions, technological and otherwise, to help them identify bots and mitigate attacks. Prior work mainly considered technical issues related to bot detection, and, except for a few efforts [5, 6, 34], mostly ignored human aspects. Nonetheless, a deeper understanding of users (how they perceive bots, the types of users who need help, the bot types they need help with, ...) can help inform the development of effective user-centered interventions.

To this end, we conducted an online study with 297 participants recruited via a crowdsourcing platform in which we 1) gauged their

perceptions of bots and whether they believe bots can influence their behavior; 2) assessed whether they can classify bots and non-bots correctly and quantified how different factors impact their accuracy; and 3) measured their satisfaction with different OSNs and how they believe OSNs should handle bots (see §3). Among others (§4), we found that participants’ perceptions of bots captured real aspects of bots’ operation (e.g., automation) and goals (e.g., artificially boosting contents’ popularity). Nonetheless, some participants confused bots with other phenomena (e.g., online tracking and personalization) or had abstract perceptions of what bots are. A substantial portion (52.53%) of the participants believed that bots do not influence them personally, including because they believed they can accurately detect bots. However, participants could not correctly classify all the accounts we showed them (31.70% avg. error). Our analysis indicated that multiple factors impact users’ likelihood to classify accounts correctly, including their perceptions of bots (e.g., users with abstract perceptions were more likely to misclassify) and the types of accounts (e.g., fake-follower bots were more likely to be misclassified than other accounts). Lastly, the participants had varied levels of satisfaction with how OSNs governed bots—they were particularly unsatisfied with OSNs they perceived to have a significant presence of bots. Based on our findings, we make several recommendations to improve interventions (§5).

2 RELATED WORK

We now review common definitions of bots, and summarize work on bot detection and human-bot interaction.

2.1 Definitions of Social Bots

Finding a commonly accepted definition of *social bots* is surprisingly challenging. Grimme et al. tackled this difficulty and concluded that most definitions attribute a level of automation to bots, but may also include motivations (e.g., maliciousness), capabilities (e.g., content generation), and additional specificity (e.g., mimicking humans) [26]. In a meta-analysis of definitions provided in 15 recent papers [3, 4, 11, 13, 17, 18, 20, 22, 29, 34, 37, 52, 53, 56, 57], we found that researchers described bots using terminology similar to some of our participants’. For example, researchers often mentioned human imitation, automated interaction with other accounts, and automated content generation as part of their definitions. Definitions also used the term “fake” and mentioned platform abuse. In contrast to researchers, however, our participants also described bots via vague or completely inaccurate terminology (§4.1).

Definitions generally identify that bots can be benign or malicious [50]. Benign bots can provide entertainment or assist humans in various means, such as aggregating news, moderating forums, and supporting customers [22, 28, 33, 34]. In contrast, malicious bots can be used in different ways, including artificially boosting influence on OSNs via “fake followers” [15], influencing people’s opinions and increasing polarization [52], spreading malware and disinformation [19, 32], and participating in spam and phishing [59].

2.2 Social Bot Detection

A significant amount of work has gone into developing bot-detection algorithms on OSNs, and especially Twitter, due to the public visibility of most of its content and a well-documented API [11, 17, 18, 37, 57]. There have also been efforts to detect bots on other OSNs, including Facebook [45] and Instagram [47].

Most bot-detection algorithms rely on machine-learning (ML) algorithms [4, 14, 18, 57]. Typically, ML models are trained to distinguish bots from non-bots using a collection of features characterizing accounts' metadata, contents of posts and profiles, networks' structures, and accounts' behavior over time. Similarly to ML algorithms, our study participants often relied on accounts' metadata and content features (e.g., follower-to-following ratio and the amount of hashtags in posts) to tell bots and non-bots apart (§4.2). Unlike algorithms, our participants used a limited set of features, and performed classification via simple, imprecise heuristics.

There have also been attempts to equip users with bot-detection tools to augment their understanding of whether accounts are likely to be bots. BotOMeter is a web app users can query with specific Twitter handles to be classified [57]. BotSight, a tool available as a browser extension or mobile app, inlines posts and profiles on Twitter with indicators denoting the likelihood that accounts are bots. While beneficial, these tools have small user bases (e.g., ~500 daily visitors to the BotOMeter website [56]), and are unlikely to be adopted by users who are expected to benefit the most from them (specifically, users who are unfamiliar with bots §5.2).

2.3 User-Bot Interaction

Prior work has discovered that users' susceptibility to being manipulated by political and spam bots is varied. Badawy et al. found that holding a conservative political ideology was a primary predictor for whether users would spread disinformation [6], while Luceri et al. attributed conservative bots' effectiveness at spreading disinformation to their network centrality [35]. Redmiles et al. characterized users that click on spam and found, among others, that users less active on the Internet were more susceptible than others [43]. By contrast, our study seeks to characterize OSN users' mental models of bots and measure users' ability to differentiate between bots and non-bots. Moreover, we do not limit ourselves to political or spam bots, but consider diverse account types.

Researchers have investigated users' ability to detect bots and disinformation. Alarifi et al. asked nine study participants to label accounts as bots and non-bots and measured 96% accuracy after classifying accounts by a vote over participants' labels [2]. Similarly, Cresci et al. asked study participants to label accounts which they later classified by a vote [16]. Consistently with our findings, many of their participants thought that bots that belong to a particular family (specifically, spam bots) were non-bots. Unlike ours, however, prior studies did not quantify the individuals' accuracy, considered limited account types, and did not explore the strategies participants used to classify accounts. Freeman concluded that the average user is not capable to accurately detect fraudulent activity, but did not explore users' perceptions and classification strategies [23]. Geeng et al. reported that their study participants—unaware of being tested—were unlikely to investigate disinformation planted into their Facebook and Twitter feeds [24]. They attributed some

participants' lack of investigation to over-confidence in their ability to spot misinformation. Our participants exhibited similar behavior (§4.1). Finally, Appling and Briscoe asked their study participants to select the features of Twitter accounts that were most associated with bots [5]. The participants denoted that links in the profile and large numbers of posts were most indicative of bots. Our participants relied on a richer feature set (§4.2).

Prior work has touched on user mental models of bots. Chaves et al. and Zamora studied user expectations from chatbots [10, 58]. They found users were dissatisfied with bots' ability to understand questions and produce quality responses. Clément et al. studied reactions to bots on Wikipedia by analyzing posts left on discussion pages [12]. They grouped the bots into two ideotypes: "servant bots" and "policing bots." Users had differing reactions to bots depending on ideotypes, with policing bots eliciting more polarized responses. Wessel et al. looked at bots as assistants during software development [55]. They surveyed developers to gauge their perceptions of bots and found the bots were not sufficiently advanced to impact the development process. Gero et al. asked participants play an online game against an automated agent and found that those who won had more accurate models of the agent's capabilities [25]. Finally, Long et al. conducted a field study on Reddit to explore users' perceptions of benign bots [34]. They found that users' perceptions varied with technical ability: while most users understood the basic bot functions, non-technical users often had unrealistic expectations from bots. Unlike prior work, we study users' perceptions of both benign and malicious bots, and do not restrict ourselves to a single OSN or automated agent when exploring perceptions. Moreover, we discover new concepts that OSN users may identify with bots, sometimes erroneously (e.g., online tracking).

3 METHODOLOGY

Here we describe our study design, analysis methods, our participants, and our study's limitations.

3.1 Study Design

To answer our research questions, we designed an online (crowd-sourced) study to survey a large and diverse group of participants. The study began by asking participants about their general OSN usage habits. In particular, we inquired about whether participants have accounts on any of the following seven OSNs: Facebook, Instagram, LinkedIn, Reddit, TikTok, Twitter, and YouTube. Then, for each OSN a participant has an account on, we asked about how often they visit the OSN, how much time they spend on it, and what content types they create or consume. We selected these specific seven OSNs because they represent a variety of popular OSNs known to host bot activity and differ in several interesting aspects, including their purpose, user-base demographics, and the modes of interactions between users. For example, TikTok is a four-year-old OSN primarily used by young users for entertainment purposes. In contrast, LinkedIn is a 17-years-old OSN used by working professionals to share professional content.

Next, we asked questions to elicit participants' perceptions of bots, and their belief on whether bots influence OSN users. Specifically, we asked participants whether they were familiar with bots and instructed them to provide concise definitions of bots and their

goals as free-form responses. If participants had not encountered the term bot before, we encouraged them to best guess the definitions and goals. Furthermore, we asked the participants how common they believe bots are on the OSNs they were active on, and whether they believe that bots influence their or others’ behavior and why. To motivate participants to submit responses that best represent their perceptions, we emphasized that the study does not aim to test them and reminded them to answer honestly.

The next part of the study aimed to gauge how accurately OSN users can differentiate between bots and non-bots. To this end, we presented the participants with 20 carefully chosen Twitter accounts and asked them to classify each account as a bot or a non-bot. This part of the study consisted of two stages. In each stage, we asked the participants to classify ten accounts, split equally between bots and non-bots. Between the two stages, we showed the participants a standard, concise definition (inspired by Grimme et al. [26] and Ferrara et al. [22]) describing bots and their goals, as a means to help inform participants who are unfamiliar with bots, and measure whether informing participants about bots improves their classification accuracy.

We decided to ask participants to classify Twitter accounts rather than accounts from other OSNs since Twitter bots are well-studied, and there are numerous public datasets we could select accounts from. We selected accounts for our study from the 14 datasets used by Yang et al. [57]. We randomly ordered the accounts available in all datasets and manually inspected individual accounts one by one to select active ones pertaining to different types identified in prior work [18, 26, 34, 50]. Specifically, we selected a diverse set of bot and non-bot account-types (benign bot, political bot, ...) differing in a variety of features, such as their popularity or the type of content they post, to measure how account characteristics affect participants’ classification. Overall, we selected 20 accounts: two groups of ten accounts each, split equally between bots and non-bots of varied types. Tab. 1 presents the selected accounts and their characteristics.

We presented participants with the accounts of one group in the first stage, followed by the accounts of the other group in the second stage. To mitigate ordering effects, we alternated between presenting group A first followed by group B and vice versa between stages, and randomized the order of accounts within each stage. Each classification question contained a screenshot of the profile page of the account (containing picture, description, number of followers, ...), followed by as many of the account’s most recent tweets as would fit on a 1,920×1,080 screen (a minimum of one tweet, and a median of 2), see Fig. 1.



Figure 1: A screenshot of a bot shown to participants in the study.

as would fit on a 1,920×1,080 screen (a minimum of one tweet, and a median of 2), see Fig. 1.

Cls	ID: Def	Group A	Group B
Non bots	Org: Organization/business	UHCF_CT	LGRforCollege
	Par: Parody, self-identifies	notzuckerberg	Queen_UK
	Pop: >1,000 followers, unverified	aaroni	GynoStar
	Upo: <1,000 followers, unverified	kb_kinsey	bablackwood
	Ver: Of public interest, verified	DineshDSouza	TheBriannaDavis
Social bots	Ben: Entertains or serves users	EndlessJeopardy	TwoHeadlines
	DSc: Publishes dating scams	brown_gabathaa	drrajashekar170
	Fak: Inflates followers	aurelias_alaura	navash73
	Pol: Backs pol. agenda, attacks others	HansonMAGOP	shinziNET
	TSc: Publishes tech scams	littlebilly_	lcj18344

Table 1: The accounts we asked participants to classify as bots and non-bots. The two account groups are listed on the right. The leftmost columns list the accounts’ classes, and type identifiers and definitions. Following best practices, we reported the malicious bots (last four rows) to Twitter.

To conclude the study, we asked participants to share their level of satisfaction with how the OSNs they have accounts on deal with bots (five-point Likert scale), and asked them how OSNs should best use bot-detection technology to moderate bots. Finally, we collected demographic information.

Our study protocol (see the supplementary material) was reviewed and approved by our organization’s legal department.

3.2 Analysis

We qualitatively analyzed participants’ definitions of bots and their goals—specifically, via inductive coding [44]—to understand participants’ mental models. Over two sessions, two coders analyzed the responses of a subset of participants to agree upon a codebook. The coders later coded the responses of a subset of additional participants independently to assess the inter-coder agreement. Upon reaching substantial agreement, the coders split the coding effort of the remaining participants’ responses equally among themselves.

Subsequently, we built logistic regression models to estimate how various factors (mental models, account types, familiarity with Twitter, etc.) impacted participants’ likelihood (i.e., odds ratio) of classifying accounts correctly. More precisely, because the responses to certain account-classification questions were *not* independent (each participant classified 20 accounts), we built mixed-effects models [46]—a specific family of models that can handle correlations between data points. For each explanatory variable, mixed-effects logistic regression models estimate how much an increment of one in explanatory variables affects the expected odds of correct account classification compared to a pre-specified baseline.

3.3 Participants

To determine the number of participants needed for the study, we ran a pilot with a convenience sample ($N=29$). We then conducted a power analysis ($\alpha=0.05$, power=0.80) and found that ≥ 288 participants were needed to confirm various factors’ impact on classification accuracy.

We administered the study via Qualtrics and recruited participants via Prolific—an online crowdsourcing platform that enables recruiting participants worldwide. To avoid self-selection bias, we did not advertise the study as one about bots. Instead, to remain inclusive, we advertised the study as one aiming to learn about experiences on different OSNs and explore usage habits. Compared to other crowdsourcing platforms (e.g., Amazon’s Mechanical Turk),

Prolific’s participants are known to produce data of higher quality [39]. To further enhance the data’s quality, we used Prolific’s pre-screening options to limit participation to workers fluent in English (the study’s language) whose submission approval-rates are higher than 95% [40]. Furthermore, we limited participation to 18-year-old or older workers who own an account on at least one of the seven OSNs our study asked about. A total of 298 participants who matched these criteria completed the study. We excluded one participant due to submitting responses irrelevant to the study’s subject matter, thus remaining with 297 participants for the analysis. Participants took an average of 22 minutes to complete the study and were compensated \$2.60 for their time.

Our participants were most active on Facebook, Instagram, and YouTube. More than 86.19% of the participants reported owning accounts on at least one of these OSNs. A substantial portion of participants reported owning accounts on Twitter (62.96%), LinkedIn (37.71%), or Reddit (53.20%). TikTok was the least popular OSN among the participants—29.29% reported owning accounts. Our participants were skewed toward younger and male populations. Participants’ ages ranged between 18 and 61, with a median of 23. Males constituted 65.32% of the total participants (34.01% reported female, and the remainder chose other). The participants were educated (29.63% college students, 44.11% owned an associate’s degree or higher) and mostly self-reported being technology savvy (76.43%). They lived in different parts of the globe, but mainly Europe (82.83%), North America (11.11%), and South America (4.71%).

3.4 Limitations

Certain limitations should be considered when interpreting our results. The account-classification task in our study is Twitter-centric. Thus, results may not precisely reflect how well users can differentiate between bots and non-bots on other OSNs. Nonetheless, we attempted to interpret our findings generally and believe that our recommendations (§5) apply to all major OSNs. We also stress that our findings about users’ (mis)perceptions are OSN-agnostic. We further note some participants (37.04%) in our study reported not owning Twitter accounts, which could have disadvantaged them in the classification task. We controlled for lacking Twitter accounts in our analysis (§4.2) and found it had no significant impact on participants’ classification accuracy.

Unlike participants in our study, OSN users may not consciously seek to differentiate between bots and non-bots. Thus, our estimate for how well users detect bots in practice may be imprecise. Still, we believe that our estimate is useful, as it likely sets an upper bound on how well OSN users perform in reality. Said differently, we expect that OSN users who do not consciously attempt to detect bots are likely to detect them less accurately than our participants.

We asked participants to classify accounts based on a screenshot of the profile page; thus not faithfully imitating a real browsing setting. Yet, we believe that our study design captures many occasions in which users consume content on OSNs via cursory examination without gathering additional information about accounts [24].

Lastly, as is common in online self-reported surveys, ours is not free of biases [27, 30]. Primarily, participants’ responses (e.g., of OSN use) may be affected by recall or desirability biases, or by misinterpreting questions. Additionally, our participants do not fully

represent the general OSN user base—they were skewed toward the young, educated, and male population. We mitigated biases via careful study design and piloting, and reminding respondents to answer thoroughly and honestly. Furthermore, to validate that our findings hold despite potential biases, we repeated our study with a different participant pool and reached consistent conclusions (§4.4).

4 RESULTS

Now we turn to our results. We start with participants’ perceptions of social bots and how they influence users. We then report participants’ account-classification accuracy and discuss the factors that impacted it. Next, we turn to participants’ estimates for bot prevalence, and discuss how these relate to participants’ satisfaction with how OSNs govern bots. Finally, we report validation results with an additional cohort of participants.

4.1 Mental Models

What are social bots? Most participants (79.80%) self-reported being familiar with the notion of bots. We asked those participants to define bots and list their goals, while we encouraged the others to make best guesses. The responses were instructive and reflected varied mental models about bots; we leveraged inductive coding to characterize them. Initially, two coders developed a codebook over two coding sessions. After analyzing 40 responses, they concluded the initial open coding process, as new themes stopped emerging. The coders then independently coded additional 40 responses to calculate the agreement among them, reaching substantial agreement (Cohen $\kappa=0.76$) [36]. Subsequently, the coders equally split the remaining responses between them for coding. Tab. 2 presents the final codebook along with example phrases, representative quotes, and the percentages of participants for which codes apply.

Participants mostly mentioned several properties describing bots. In most cases, the definitions captured actual properties of bots. A large portion of the participants (49.83%) indicated automation as part of their definition, using terms such as artificial intelligence and algorithm to describe how bots operate. A substantial chunk of participants (20.20%) described bots as “fake” or ungenue accounts without specifying how they are controlled, often referring to so-called fake followers that users can buy to promote the their accounts’ popularity [47]. Some even revealed that they or someone they know have once purchased fake followers. Various participants (19.19%) described bots as interactive accounts that communicate or interact with other accounts, while some (16.50%) noted that bots often pretend to be humans. A few participants (3.70%) suggested that bots use specific technologies to produce original content, and a smaller set (1.01%) described bots as anonymous users. Other participants provided vague or inaccurate definitions. A substantial amount of participants defined bots in abstract terms (17.51%), describing them as “something” that performs specific tasks, while certain participants provided cyclical definitions (12.79%). Some participants (6.40%) confused bots with other technologies, primarily cookies and other tracking technologies, and a small number of participants (2.02%) thought bots were sock puppets—users that masquerade as others for deception [8]. Lastly, a few participants stated they do not know what bots are (2.36%) or provided unintelligible responses (1.68%).

Cat.	Code	%	Example(s)	Quote
Definition	Automated	49.83%	AI; agent; program	P145: "A form of AI that contributes to a discussion."
	Fake	20.20%	"Fake"; ungenune	P137: "It's a fake account that can create content or interact with other users."
	Interactive	19.19%	Communicate; interact	P170: "An automatic account that interacts with other accounts, for example makes comments."
	Abstract	17.51%	"Something"	P139: "A bot is something that posts or interacts with other people online."
	Impersonate	16.50%	Pretend to be human	P157: "[...] pretends to be a user, mimicking human behavior on Internet."
	Cyclical	12.79%	Social bot; bot	P198: "It is an account that is managed by a bot..."
	Confuse	6.40%	Cookie; trackers	P37: "Are they like clever cookies that influence the content you see?"
	Create content	3.70%	Create original content	P284: "It is designed to create posts. Mainly used as automated replies."
	Don't Know	2.36%	-	P149: "I don't know what is social bot..."
	Sock puppet	2.02%	Masquerade as others	P121: "[...] humans [...] running many accounts [...] with the purpose of persuading other genuine users."
	Unintelligible	1.68%	-	P281: "Social bot is not completed project nowadays."
	Anonymous	1.01%	Anonymous user	P271: "...an anonymous 'user,' which [...] boosts up the popularity of certain accounts, persons, politicians."
	Goal	Boost	42.09%	Inflate popularity
Influence		31.31%	Sway public opinion	P174: "To influence your opinions."
Benign		11.11%	Help users	P26: "[...] communicative agent that can be called in the conversation to give you any particular info."
Confuse		10.44%	Personalize; recommend	P167: "Making things happen automatically by using certain algorithms, like choosing who will get which ads."
Scam		9.76%	Spam; commit crime	P156: "Some try to trick you in clicking on malicious links."
Advertise		9.09%	Promote products	P264: "[...] to interact with people with some content for e.g with some advertisement about some product."
Earn money		8.08%	Make money	P215: "[Bots' goals are] Political influence or monetary gain."
Support		6.40%	Customer support	P12: "Make the customer service work easier for companies..."
Don't Know		4.71%	-	P242: "Sincerely, I don't know the goals of such social bots."
Unintelligible		4.04%	-	P281: "I believe social bot is a good friend of all people."

Table 2: Codes describing mental models of bots. The second column from the left reports the percentage of participants for which codes apply; the third column lists examples of common words and phrases used by participants; and the rightmost column provides representative participant quotes. Codes highlighted in gray describe participants' responses, not bots.




Figure 2: Participants believed that their own behavior was not influenced by bots, but that others' was.

The participants articulated varied goals, reflecting the real-life diversity in the usage of bots. Multiple participants (42.09%) suggested that bots sought to artificially boost or inflate the popularity of accounts or posts. Many (31.31%) stated that bots aim to influence users or sway public opinion (e.g., via disinformation). Some (11.11%) thought of bots as benign actors providing features to help users—ranging from entertainment to automated question answering—with a subset specifically citing customer support (6.40%) and legitimate advertising (9.09%). Participants also identified that bots can be used to scam OSN users and spread spam (9.76%), and are often operated to earn their operators monetary gains (8.08%). A substantial group (9.76%) confused bots' goals with those of other technologies, such as recommending content to users and personalizing the user experience. Lastly, a few participants were unable to articulate the goals of bots (4.71%) or submitted unintelligible responses (4.04%).

Who is influenced by bots? We asked participants whether they thought bots influenced their or others' behavior on OSNs. Fig. 2 presents the results. Most participants (52.53%) thought that bots were unlikely to influence them. To explain why, participants mentioned that they do not interact with bots, do not believe everything they read, are challenging to mislead, and can spot bots accurately. This is consistent with Geeng et al.'s finding that most people do not believe they are susceptible to propaganda and disinformation [24]. Somewhat paradoxically, however, participants

misclassified a significant portion of bots as non-bots (§4.2), indicating that bots are likely to unconsciously deceive them in reality.

By contrast, most participants (62.63%) believed that were likely influenced by bots. Among others, the participants explained their belief by noting that others likely do not know what bots are, cannot distinguish bots from non-bots, and are likely to be misled into thinking that certain content is more popular than it truly is.

4.2 Users' Ability to Detect Bots

Accuracy. We examined how well participants were able to distinguish between bots and non-bots. Overall, participants classified 68.30% ($\pm 12.68\%$) of the accounts correctly. Namely, participants correctly classified about 14 of the 20 accounts on average, with most participants (76.43%) making between 11 and 17 correct classifications. At the same time, none of the participants could classify all 20 accounts correctly, with a substantial amount of participants (10.10%) misclassifying $\geq 50\%$ of the accounts.

When examining misclassifications, we found that participants were more likely to misclassify bots as non-bots than the other way around. More specifically, participants misclassified an average of 23.27% of non-bots as bots, and 40.13% of bots as non-bots. This is worrisome, because by tricking participants into believing that they are non-bots, malicious bot owners are more likely to accomplish their goals (e.g., carry out a successful scam).

Factors. We built a mixed-effects regression model to assess how various factors impacted the participants' likelihood of classifying accounts correctly. The dependent variable in the model was an indicator for whether a participant answered a certain classification question correctly (1 for a correct answer, 0 otherwise). We derived the explanatory (a.k.a. independent) variables from the survey responses and our qualitative analysis. We used a random intercept as a random effect in the model to account for differences between participants' baseline likelihoods to classify correctly.

Category	Variable	Odds	Conf. int.	p-val.
Intercept	Mean	3.10	[2.44, 3.94]	<0.01
	Std	0.30	-	-
Demog.	AgeScaled	0.63	[0.39, 0.99]	0.04
	HighDegree	1.31	[1.14, 1.52]	<0.01
Habits	HeavyReddit	1.38	[1.17, 1.62]	<0.01
Percep.	DefAbstract	0.76	[0.63, 0.94]	0.01
	DefAuto	0.84	[0.71, 0.98]	0.03
	DefImpers.	1.30	[1.06, 1.57]	0.01
	GoalDntKw	0.67	[0.49, 0.91]	0.01
	GoalScam	1.43	[1.13, 1.82]	<0.01
	GoalUnint.	0.65	[0.47, 0.91]	0.01
Acc. Type.	BotBen	0.66	[0.51, 0.87]	<0.01
	BotDsc	0.84	[0.64, 1.08]	0.17
	BotFak	0.19	[0.15, 0.24]	<0.01
	BotPol	0.47	[0.34, 0.64]	<0.01
	BotTsc	0.58	[0.45, 0.75]	<0.01
	NonBotPar	0.58	[0.45, 0.75]	<0.01
	NonBotPop	2.44	[1.79, 3.29]	<0.01
	NonBotUpo	1.51	[1.14, 1.99]	<0.01
	NonBotVer	1.35	[1.03, 1.77]	0.03
	Stage	SawDef	0.84	[0.74, 0.97]
SawDef:BotPol		1.52	[1.06, 2.18]	0.02

Table 3: Logistic regression model’s parameter estimates after model selection. For each variable we provide the estimated mean odds, the 95% confidence intervals, and the p -value for the null hypothesis that the mean odds are equal to 1. Except for BotDsc, all mean odds are estimated to be different than 1 with statistical significance (p -value<0.05).

Initially, we incorporated a rich set of independent variables in our model, including ones describing 1) the OSNs that participants were heavily active on (e.g., spending at least an hour a day); 2) whether participants indicated being familiar with bots; 3) the study completion time; 4) the account type being classified (dating scam bot, verified non-bot, ...); 4) participants’ demographics; 5) participants’ mental models of bots and their goals; 6) whether participant saw the bot definition we provided (i.e., whether they were in the first or second stage of the classification task); 7) various interactions between whether the participant saw the definition and other variables (account type, familiarity with bots, ...). The initial model and variables are available in the supplementary material. Next, we followed a standard backward model-selection process to simplify the model while maintaining a good fit [46]. Specifically, we gradually removed the variables with the highest p -values, one at a time, until reaching a point at which the Bayesian Information Criterion (BIC) decreased by <5. Eventually, the process yielded the model with the variables listed in Tab. 3 along with their estimates ($R^2=15.81\%$, $BIC=7010$). The model explains 15.81% of the variance in the data, indicating that there are missing factors that can further explain participants’ classification accuracy. Nonetheless, in a 5-fold cross-validation process, we found that a predictive model using only the fixed effects can predict when participants classified accounts correctly or not with high accuracy (70.77% mean accuracy and 68.16% Receiver-Operating Characteristic area under curve).

The baseline in the model is an 18-years-old user with a high-school or lower education level, who does not use Reddit heavily, and who is classifying an organizational non-bot account (Org from Tab. 1) before seeing the definition. From the intercept, we can see that the mean odds that the baseline user will classify the account correctly is 3.1. Said differently, the mean probability of correct

classification is $\frac{3.1}{3.1+1} \approx 0.76$. Compared to the baseline, a user of the maximal age we encountered in our study (61) is estimated to have $\times 0.62$ lower mean odds of correct classification (every year of age decreases the odds by about $\times 0.99$), while a user with an associate’s degree or higher is likely to have $\times 1.31$ higher mean odds of correct classification. Other demographic properties (e.g., gender) and the study-completion time had no statistically significant correlation with classification’s correctness (i.e., these variables were dropped by model selection).

The types of OSNs that participants reported using heavily had little impact on their classification accuracy—except for the variable indicating heavy Reddit usage, none of the others survived model selection. Interestingly, despite asking participants to classify Twitter accounts, heavy Twitter users did not perform better than others. Heavy Reddit users, however, had $\times 1.38$ higher mean odds of correct classification than the baseline. We hypothesize that this is due to Reddit users being particularly young and educated [41] and their constant exposure to benign bots [34].

The model also tells us that the users’ perceptions of bots impact their classification accuracy. Notably, participants who provided abstract definitions of bots and ones who mentioned not being aware of bots’ goals had $\times 0.67$ – $\times 0.76$ lower estimated mean odds for correct classification than the baseline. In contrast, participants who mentioned that bots may attempt to impersonate humans and mentioned scams as a potential goal of bots had $\times 1.30$ – $\times 1.43$ higher estimated mean odds than the baseline. This indicates that helping users form more accurate mental models of bots may significantly improve their ability to distinguish between bots and non-bots.

We can further learn from the model that the types of accounts being classified impacted participants’ accuracy. Except for bots used for dating scams, participants’ estimated mean odds of classifying bots correctly was $\times 0.19$ – $\times 0.66$ lower than the baseline, with fake followers being particularly difficult to classify correctly. By contrast, non-bots were generally easier for participants to classify correctly—participants exhibited $\times 1.35$ – $\times 2.44$ higher estimated mean odds to classify unpopular, popular, and verified non-bots correctly compared to the baseline. The exceptions are parody non-bots with significantly lower estimated mean odds ($\times 0.58$) of correct classification compared to the baseline. These findings highlight that aiding users in identifying bots via specialized interventions targeting specific types of accounts could be a promising avenue.

Finally, we can learn from Tab. 3 that seeing the definition before the second stage of account classification did not improve the participants’ performance: while the estimated mean odds of correctly classifying political bots increased by $\times 1.28$ after participants saw the definition, the odds decreased by $\times 0.84$ in all other cases. Potential explanations are that the specific definition we provided was not sufficiently accessible to the participants, repeated exposures to definitions might be needed to improve participants’ accuracy, or that participants paid slightly less attention to questions later in the study. Thus, further research is needed to find out how to develop an accessible definition of bots to help users identify such accounts in practice and determine the definition’s effectiveness.

Classification heuristics. Participants relied on various rules of thumb to detect bots. Certain participants based their classifications on simple rough statistics from accounts’ *metadata*, such as the number of tweets and followers, the ratio between the number

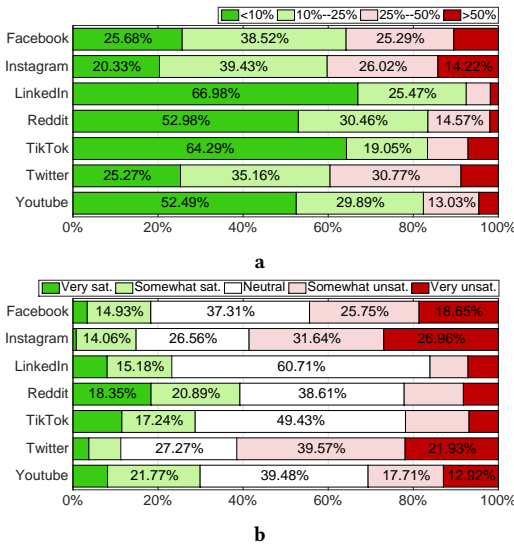


Figure 3: (a) Participants' perceptions of bots' prevalence (percentage of accounts) on OSNs. (b) Participants' satisfaction with bots' governing.

of accounts followed and the number of followers, the amount of retweeting of others' posts, and the number of hashtags in tweets. Other participants mainly relied on the *contents* of tweets and profile pages, including whether they contain suspicious links, whether content included promotions and marketing material, the existence of grammatical and spelling mistakes in posts, and repetitiveness and originality of posts (e.g., whether they repeated titles of linked articles). Lastly, other participants relied on *appearance and feel* and stated that bots' tweets seemed "unnatural" and "strange," their photos appeared seductive or unrealistic (e.g., in dating scams), and that their posts seemed "sketchy." The rough statistics used by participants were similar to those used by ML algorithms for bot detection. Nonetheless, participants' heuristics were generally less exhaustive, precise, and systematic than ML algorithms, thus helping explain participants' relatively high misclassification rates.

4.3 Perceptions of Bot Prevalence

When asked to estimate the prevalence of bots on the OSNs they own accounts on, participants believed they were most prevalent on Facebook, Instagram, and Twitter (Fig. 3a). Many participants ($\geq 35.80\%$) expected that $\geq 25\%$ of the accounts on those OSNs are bots. While participants likely overestimated the actual prevalence of bots—Varol et al. estimated that 9%–15% of active accounts on Twitter were bots [53]—their responses reflect that bots on certain OSNs are likely responsible for a substantial amount of the interactions with humans.

Participants expressed higher satisfaction levels with how OSNs governed bots when they perceived bots were less prevalent on the OSNs, and vice versa (Fig. 3b). Specifically, participants were less satisfied with how Facebook, Instagram, and Twitter dealt with bots than with the other OSNs. Unsatisfied participants attributed their lack of satisfaction to the lack of OSNs' efforts to moderate bots (e.g., by removing inappropriate adult content or propaganda), which they sometimes interpreted as lack of care for users; being

frequently bothered by automated comments, direct messages, and connection requests; and preference to interact with humans. On the other hand, neutral participants expressed that it was unavoidable to have some bots on the platform; detecting bots should not entirely be OSNs' job, as users should develop a healthy skepticism toward the content they consume; and they do not encounter bots often. Finally, satisfied participants attributed their satisfaction to OSNs' efforts to moderate bots which they felt were tangible (e.g., OSNs swiftly removing content flagged as spam) and clearly publicized; the utility they found in (benign) bots; and to not being exposed to scams.

4.4 Additional Validation

Because our main set of participants was skewed, we recruited another cohort of participants to validate our findings' generalizability. Specifically, we recruited 139 volunteers from our organization—a global software company—to participate in the study. This validation sample was more balanced with respect to age (40 vs. 23 median age) and gender (55.40% vs. 65.32% male) than the participants recruited via Prolific. The validation participants resided mostly in North America (65.47%), Europe (20.14%), and Asia (13.67%). As expected, a large portion (37.41%) identified as computer engineers, but many identified as finance, legal, and administrative support professionals, among others. TikTok and Reddit ($\leq 19.42\%$ owned accounts) were the least popular OSNs among the participants. Most participants (59.71%) owned Twitter accounts, and $\geq 64.03\%$ owned accounts on at least one of the remaining OSNs.

The validation participants exhibited similar mental models of bots as the Prolific participants and obtained equivalent account-classification accuracy (67.37%). Compared to the Prolific participants, the validation participants erred slightly more when classifying non-bots (31.29% misclassification), but did better on bots (33.96% misclassification). The relatively small sample size of the validation dataset precluded fitting a regression model on the dataset itself. Therefore, to test whether the validation data yield parameter estimates similar to the Prolific data, we fit a mixed-effects logistic regression containing all parameters listed in Tab. 3, while adding an indicator variable for the datasets and an interaction term between the dataset and each of the original parameters. The resulting model helped us estimate whether the validation data yield different slopes or intercept compared to the original data. We found no statistically significant difference between the estimates of the intercepts and 15 of the 20 parameters across the two datasets, while three of the five remaining parameters had a similar correlation type (i.e., positive or negative) with the dependent variable for the two datasets. To further increase our confidence in the findings, we tested how well can a model trained on the Prolific data using only fixed effects predict whether the validation participants accurately classified accounts. The model achieved 71.26% accuracy and 57.80% Receiver-Operating Characteristic area under curve—comparable to the performance achieved on the Prolific data. Finally, the validation participants' satisfaction with how the different OSNs dealt with bots was similar to that of the Prolific participants. Primarily, the participants were least satisfied with Facebook, Instagram and, Twitter. Thus, overall, conclude that our findings on the Prolific data generalize well to our validation data.

5 DISCUSSION

Our study demonstrates that OSN users struggle in differentiating between bots and non-bots. Even worse, users are especially likely to mistake malicious bots as non-bots, thus rendering them susceptible to disinformation and attacks. Hence, there is a dire need for interventions to help users identify different OSN account-types. We now discuss potential technical and education-based interventions.

5.1 User Education

Our study shows that participants with abstract perceptions of bots and those unaware of bots' goals were more likely than others to misclassify accounts. In contrast, participants who were familiar with the specificities of bots (e.g., usage to spread scams) detected bots with higher accuracy than others. These results indicate that user education and habituation to bots can significantly boost their ability to distinguish between bots and non-bots. We believe that education-based interventions can be useful, but they may be insufficient to prevent account misclassification, as even expert users who were familiar with bots misclassified accounts.

5.2 Systems and Tools

ML-based bot-detection algorithms differentiate between bots and non-bots with significantly higher accuracy than our participants. Using the same type of information that was available to our participants, algorithms often detect >95% of bots while misclassifying a small portion (<5%) of non-bots [31, 57]. In comparison, our participants misclassified 40.13% of the bots and 23.27% of the non-bots we asked about. The participants' lower accuracy can be attributed to using less comprehensive and systematic classification heuristics than algorithms. Additionally, bots can be abnormal in non-obvious ways; for instance, high follower growth-rates are useful features in bot detection [57]. Computing this and similar metrics is possible via OSN-provided data, but doing so manually for many accounts is infeasible. OSNs can help users detect bots by making such metrics visible, thus expanding the limited feature set they typically rely on. More generally, ML-based algorithms can empower users to detect bots accurately while avoiding disinformation and attacks.

Where to deploy detection? Detection algorithms can be deployed at various locations. One possibility is to make them available via OSN-independent websites, such as BotOMeter [56], that users would visit to classify certain accounts, or browser extensions, such as BotSight [38], that annotate accounts as bots and non-bots in situ, as users browse OSNs. This deployment model suffers from a scalability issue. Because users need to visit dedicated websites or install extensions actively, the user base of OSN-independent tools tends to be relatively small, thus leading to a limited impact. For example, out of the hundreds of millions of OSN users, the BotOMeter website had ~500 daily visits as of two years ago [56]. We also expect that the users who are likely to benefit the most from detection algorithms—those who are less familiar with bots or falsely believe that they can accurately detect bots—are unlikely to use OSN-independent tools.

Another, preferable, possibility is to deploy detection algorithms internally, behind the scenes, on OSN platforms. Unlike in the previous deployment model, all OSN users would benefit from the

detection results. Indeed, certain OSNs like Reddit rely heavily on bot-detection systems to detect and suspend malicious bots [42]. We found that users of such OSNs are more satisfied with how bots are dealt with compared to other OSNs that are known to have a more lax policy towards bots.

How to use detection results? Once they detect bots, systems and tools might react in a variety of ways. For example, they might suspend the accounts [49] or annotate their posts as automatically produced or erroneous [7, 38]. To better understand users' needs, we asked our participants to choose the course of action they would prefer OSNs take once they detect bots. For benign and malicious bots, participants could choose between suspending accounts, adding indicators saying accounts are bots, or taking no action whatsoever. While systems and tools can conceivably perform a broader set of actions (e.g., move bot-produced content to a designated OSN section), we hoped that participants' responses would shed light on whether there is a strong preference toward a specific course of action or whether different users might have different preferences for how OSNs should handle bots.

Our participants' responses were heterogeneous. Most participants (61.62%) stated that they would like detection tools to suspend malicious bots and allow benign bots while showing clear indicators identifying them as such. A large portion (21.21%) expressed that they would prefer OSNs to suspend all bots, regardless of whether they are malicious or benign, and another substantial portion (11.78%) expressed that they would prefer OSNs to allow all bots while showing indicators. The lack of consensus between participants indicates that OSNs would better serve users by providing them with controls over how different bot types are treated. Currently available systems and tools, however, typically cannot distinguish between malicious and benign bots, and are programmed to take a single predetermined action—e.g., adding indicators on all bots [38]—that users cannot adjust. Additionally, bot-detection algorithms are typically evaluated on various accounts coarsely categorized as bots and non-bots, regardless of the specific type. These algorithms can potentially support users better if they help users detect specific account types that they are likely to misclassify (e.g., fake followers) with higher accuracy.

6 CONCLUSION

We studied 297 OSN users to improve our understanding of how they perceive social bots and measure their ability to detect them. While we found that participants' mental models often capture fundamental aspects of bots, some participants did not know what bots are, had an abstract understanding of bots, or confused them with other phenomena. Our study highlights participants' difficulty distinguishing between bots and non-bots: when asked to classify 20 accounts, all participants misclassified one or more accounts, with 10.10% of the participants misclassifying at least ten accounts. Mental models (e.g., abstract understanding of bots) and types of accounts being classified (e.g., fake followers) played central roles in participants' ability to classify accounts correctly. The participants showed a lack of satisfaction with how certain OSNs govern bots and expressed a desire for varied intervention types. Our findings help inform us about the benefits of educational and technical interventions and ways to improve them.

REFERENCES

- [1] Norah Abokhodair, Daisy Yoo, and David W McDonald. 2015. Dissecting a social botnet: Growth, content and influence in Twitter. In *CSCW*.
- [2] Abdulrahman Alarifi, Mansour Alsaleh, and AbdulMalik Al-Salman. 2016. Twitter turing test: Identifying social machines. *Information Sciences* 372 (2016).
- [3] Seyed Ali Alhosseini, Raad Bin Tareaf, Pejman Najafi, and Christoph Meinel. 2019. Detect me if you can: Spam bot detection using inductive representation learning. In *WWW Companion*.
- [4] Eiman Alothali, Nazar Zaki, Elfadil A Mohamed, and Hany Alashwal. 2018. Detecting social bots on Twitter: A literature review. In *IIT*.
- [5] Darren Scott Appling and Erica J Briscoe. 2017. The perception of social bots by human and machine. In *IFC*.
- [6] Adam Badawy, Kristina Lerman, and Emilio Ferrara. 2019. Who falls for online political manipulation?. In *WWW Companion*.
- [7] BBC. 2020. Twitter tags Trump tweet with fact-checking warning. <https://bbc.in/34Weicy>. Last accessed on 06-19-2022.
- [8] Zhan Bu, Zhengyou Xia, and Jiandong Wang. 2013. A sock-puppet detection algorithm on virtual spaces. *Knowledge-Based Systems* 37 (2013), 366–377.
- [9] BuzzFeed. 2020. Whistleblower says Facebook ignored global political manipulation. <https://bit.ly/33zB0GB>. Last accessed on 06-19-2022.
- [10] Ana Paula Chaves and Marco Aurelio Gerosa. 2021. How should my chatbot interact? A survey on human-chatbot interaction design. In *HCI*.
- [11] Nikan Chavoshi, Hossein Hamooni, and Abdullah Mueen. 2016. DeBot: Twitter bot detection via warped correlation. In *ICDM*.
- [12] Maxime Clément and Matthieu J Guittou. 2015. Interacting with bots online: Users' reactions to actions of automated programs in Wikipedia. *Computers in Human Behavior* 50 (2015), 66–75.
- [13] Stefano Cresci. 2019. Detecting malicious social bots: Story of a never-ending clash. In *MISDOOM*.
- [14] Stefano Cresci. 2020. A decade of social bot detection. *Commun. ACM* 63, 10 (2020), 72–83.
- [15] 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.
- [16] Stefano Cresci, Roberto Di Pietro, Marinella Petrocchi, Angelo Spognardi, and Maurizio Tesconi. 2017. The paradigm-shift of social spambots: Evidence, theories, and tools for the arms race. In *WWW*.
- [17] Clayton Allen Davis, Onur Varol, Emilio Ferrara, Alessandro Flammini, and Filippo Menczer. 2016. BotOrNot: A system to evaluate social bots. In *WWW Companion*. 273–274.
- [18] Juan Echeverria, Emiliano De Cristofaro, Nicolas Kourtellis, Ilias Leontiadis, Gianluca Stringhini, and Shi Zhou. 2018. LOBO: Evaluation of generalization deficiencies in Twitter bot classifiers. In *ACSAC*. 137–146.
- [19] Manuel Egele, Gianluca Stringhini, Christopher Kruegel, and Giovanni Vigna. 2013. COMPA: Detecting compromised accounts on social networks. In *NDSS*.
- [20] Emilio Ferrara. 2020. Bots, elections, and social media: A brief overview. *Disinformation, Misinformation, and Fake News in Social Media* (2020), 95–114.
- [21] Emilio Ferrara. 2020. What types of COVID-19 conspiracies are populated by Twitter bots? *First Monday* (2020).
- [22] Emilio Ferrara, Onur Varol, Clayton Davis, Filippo Menczer, and Alessandro Flammini. 2016. The rise of social bots. *Commun. ACM* 59, 7 (2016), 96–104.
- [23] David Mandell Freeman. 2017. Can you spot the fakes? On the limitations of user feedback in online social networks. In *WWW*. 1093–1102.
- [24] Christine Geeng, Savanna Yee, and Franziska Roesner. 2020. Fake news on Facebook and Twitter: Investigating how people (don't) investigate. In *CHI*.
- [25] Katy Ilonka Gero, Zahra Ashktorab, Casey Dugan, Qian Pan, James Johnson, Werner Geyer, Maria Ruiz, Sarah Miller, David R Millen, Murray Campbell, et al. 2020. Mental models of ai agents in a cooperative game setting. In *CHI*.
- [26] 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.
- [27] Andrew Guess, Kevin Munger, Jonathan Nagler, and Joshua Tucker. 2019. How accurate are survey responses on social media and politics? *Political Communication* 36, 2 (2019), 241–258.
- [28] Taylor Hatmaker. 2018. This bot unrolls Twitter threads and turns them into readable blog posts. <https://tcrn.ch/31pTsRY>. Last accessed on 06-19-2022.
- [29] Stefanie Hausteijn, Timothy D Bowman, Kim Holmberg, Andrew Tsou, Cassidy R Sugimoto, and Vincent Larivière. 2016. Tweets as impact indicators: Examining the implications of automated "bot" accounts on Twitter. *Journal of the Association for Information Science and Technology* 67, 1 (2016), 232–238.
- [30] Ruogu Kang, Stephanie Brown, Laura Dabbish, and Sara Kiesler. 2014. Privacy attitudes of mechanical turk workers and the us public. In *SOUPS*.
- [31] Sneha Kudugunta and Emilio Ferrara. 2018. Deep neural networks for bot detection. *Information Sciences* 467 (2018), 312–322.
- [32] Srijan Kumar and Neil Shah. 2018. False information on web and social media: A survey. *arXiv preprint arXiv:1804.08559* (2018).
- [33] Tetyana Lokot and Nicholas Diakopoulos. 2016. News Bots: Automating news and information dissemination on Twitter. *Digital Journalism* 4, 6 (2016), 682–699.
- [34] Kiel Long, John Vines, Selina Sutton, Phillip Brooker, Tom Feltwell, Ben Kirman, Julie Barnett, and Shaun Lawson. 2017. "Could you define that in bot terms"? Requesting, creating and using bots on Reddit. In *CHI*.
- [35] Luca Luceri, Ashok Deb, Adam Badawy, and Emilio Ferrara. 2019. Red bots do it better: Comparative analysis of social bot partisan behavior. In *WWW Companion*.
- [36] Mary L McHugh. 2012. Interrater reliability: the kappa statistic. *Biochemia medica: Biochemia medica* 22, 3 (2012), 276–282.
- [37] Amanda Minnich, Nikan Chavoshi, Danai Koutra, and Abdullah Mueen. 2017. BotWalk: Efficient adaptive exploration of Twitter bot networks. In *ASONAM*.
- [38] NLOK. 2020. BotSight. <https://bit.ly/2RHHv8c>. Last accessed on 06-19-2022.
- [39] Eyal Peer, Laura Brandimarte, Sonam Samat, and Alessandro Acquisti. 2017. Beyond the Turk: Alternative platforms for crowdsourcing behavioral research. *Journal of Experimental Social Psychology* 70 (2017), 153–163.
- [40] Eyal Peer, Joachim Vosgerau, and Alessandro Acquisti. 2014. Reputation as a sufficient condition for data quality on Amazon Mechanical Turk. *Behavior research methods* 46, 4 (2014), 1023–1031.
- [41] PRC. 2019. Who uses YouTube, WhatsApp and Reddit. *Pew Research Center* (2019).
- [42] Reddit. 2020. Reddit Security Report. <https://bit.ly/2SVJSBI>. Last accessed on 06-19-2022.
- [43] Elissa M Redmiles, Neha Chachra, and Brian Waismeyer. 2018. Examining the demand for spam: Who clicks?. In *CHI*.
- [44] Johnny Saldaña. 2015. *The coding manual for qualitative researchers*.
- [45] Giovanni C Santia, Munif Ishad Mujib, and Jake Ryland Williams. 2019. Detecting social bots on Facebook in an information veracity context. In *ICWSM*.
- [46] Howard J Seltman. 2012. Experimental design and analysis.
- [47] Indira Sen, Anupama Aggarwal, Shiven Mian, Siddharth Singh, Ponnurangam Kumaraguru, and Anwitaman Datta. 2018. Worth its weight in likes: Towards detecting fake likes on instagram. In *WebSci*.
- [48] Mohammad Shafahi, Leon Kempers, and Hamideh Afsarmanesh. 2016. Phishing through social bots on Twitter. In *Big Data*. IEEE.
- [49] Craig Timberg and Elizabeth Dwoskin. 2018. Twitter is sweeping out fake accounts like never before, putting user growth at risk. <https://wapo.st/3llM8T>. Last accessed on 06-19-2022.
- [50] Milena Tsvetkova, Ruth Garcia-Gavilanes, Luciano Floridi, and Taha Yasseri. 2017. Even good bots fight: The case of Wikipedia. *PLoS one* 12, 2 (2017).
- [51] Aman Tyagi, Matthew Babcock, Kathleen M Carley, and Douglas C Sicker. 2020. Polarizing tweets on climate change. In *SBP-BRIMS*.
- [52] Joshua Uyheng and Kathleen M Carley. 2019. Characterizing bot networks on Twitter: An empirical analysis of contentious issues in the Asia-Pacific. In *SBP-BRIMS*.
- [53] Onur Varol, Emilio Ferrara, Clayton A Davis, Filippo Menczer, and Alessandro Flammini. 2017. Online human-bot interactions: Detection, estimation, and characterization. In *ICWSM*.
- [54] Kurt Wagner and Rani Molla. 2018. Facebook has disabled almost 1.3 billion fake accounts over the past six months. <https://bit.ly/3knOsEA>. Last accessed on 06-19-2022.
- [55] Mairieli Wessel, Bruno Mendes De Souza, Igor Steinmacher, Igor S Wiese, Ivanilton Polato, Ana Paula Chaves, and Marco A Gerosa. 2018. The power of bots: Characterizing and understanding bots in OSS projects. In *HCI*.
- [56] Kai-Cheng Yang, Onur Varol, Clayton Davis, Emilio Ferrara, Alessandro Flammini, and Filippo Menczer. 2019. Arming the public with artificial intelligence to counter social bots. *Human Behavior and Emerging Tech* (2019).
- [57] Kai-Cheng Yang, Onur Varol, Pik-Mai Hui, and Filippo Menczer. 2020. Scalable and generalizable social bot detection through data selection. In *AAAI*.
- [58] Jennifer Zamora. 2017. I'm sorry, Dave, I'm afraid I can't do that: Chatbot perception and expectations. In *Proc. HAI*.
- [59] Xianchao Zhang, Shaoping Zhu, and Wenxin Liang. 2012. Detecting spam and promoting campaigns in the Twitter social network. In *ICDM*.