Journal of Computational Social Science (2022) 5:1409–1425
https://doi.org/10.1007/s42001-022-00173-9

RESEARCH ARTICLE

Comparative analysis of social bots and humans during the COVID-19 pandemic

Ho-Chun Herbert Chang[1,2] · **Emilio Ferrara**[2,3]

Received: 27 March 2022 / Accepted: 14 May 2022 / Published online: 30 June 2022
© The Author(s), under exclusive licence to Springer Nature Singapore Pte Ltd. 2022

Abstract
Using more than 4 billion tweets and labels on more than 5 million users, this paper compares the behavior of humans and bots politically and semantically during the pandemic. Results reveal liberal bots are more central than humans in general, but less important than institutional humans as the elite circle grows smaller. Conservative bots are surprisingly absent when compared to prior work on political discourse, but are better than liberal bots at eliciting replies from humans, which suggest they may be perceived as human more frequently. In terms of topic and framing, conservative humans and bots disproportionately tweet about the Bill Gates and bioweapons conspiracy, whereas the 5G conspiracy is bipartisan. Conservative humans selectively ignore mask-wearing and we observe prevalent out-group tweeting when discussing policy. We discuss and contrast how humans appear more centralized in health-related discourse as compared to political events, which suggests the importance of credibility and authenticity for public health in online information diffusion.

Keywords Social bots · Twitter · Covid-19 · Political asymmetry

Ho-Chun Herbert Chang and Emilio Ferrara contributed equally to this work.

✉ Ho-Chun Herbert Chang
herbert.hc.chang@gmail.com

Emilio Ferrara
emiliofe@usc.edu

[1] Annenberg School for Communication and Journalism, University of Southern California, Los Angeles, CA 90007, USA

[2] Information Science Institute, Los Angeles, CA 90292, USA

[3] Viterbi School of Engineering, USC, Los Angeles, CA 90007, USA

Introduction

Since 2016, when social bots were first identified as significant participant on social media, debates centered around their impact on the health of our public discourse [1]. Social media continues to grow as a significant channel for which users share information, engage in debate, and importantly, receive and consume the news. Prior studies have brought to light how social media platforms facilitated manipulation and distortion [2–4]. Social media platforms Misinformation in particular has been identified as both a symptom and cause for growing polarization within the United States [5, 6], which motivates research on identifying the actors and narratives around these issues. Accuracy, truthfulness, and authenticity in absence of misrepresentation, are crucial for a representative democracy.

The COVID-19 pandemic is the latest case-study. While the pandemic is a health crisis, it was as much a political one in that misperceptions and misinformation are often modulated by different layers of identity, including race, religion, and partisanship [7, 8]. Health misinformation regarding treatment (such as hydrochloroquine and other unverified treatments) ran rampant and aligned with political parties [9, 10]. Furthermore, misinformation narratives vary wildly across different countries, each driven by different institutional leaders, media ecosystems, levels of automation, and thus requiring different interventions [11].

As such, the ways in which automated accounts intersect with truthfulness, civic engagement, and political polarization remains a critical avenue of study. While prior studies have studied the discourse that bots promulgate online, only recently have studies targeted systematic measure of how bots engage and compare to humans in terms of the homogeneity of their consumption [12]. Additionally, a few refinements in the field have emerged. First, there has been an increase in literature that show how bots need not be always a negative force on social media. For instance, Monsted et al. showed bots can potentially encourage good behaviors, such as vaccination uptake [13]. Methodologically, a field study from Bail et al. further used them as a tool for researching online polarization and news exposure [5]. Secondly, as automation schemes becomes increasingly sophisticated, more classes of bots have appeared and respective means of identifying them [14].

Our primary objective in this paper is to assess how bots in engaging humans, the extent that they contribute to political extremism, and lingustical features unique to humans and bots. Each of these have important implications to theories of how information flow, and how partisan identity modulates this flow [15]. Leveraging a dataset of more than 4 billion tweets on the COVID-19 pandemic and labels on 5.4 million users, we investigated the interactions between bot/human identity and political diet (which we will use interchangably with partisanship). Our cross-sectional analysis thus considers (a) red bots and blue bots and (b) bots and their respective human counterparts, with the following research questions:

1. Bots versus bots:

   (a) *Are there difference in how liberal and conservative bots engage humans?* We consider differences between how bots engage with COVID-19 discussion.
   (b) *How effective are bot strategies?* We utilize three metrics introduced by Luceri et. al. [16] that quantifies the ability for bots to elicit retweets and replies.

2. Bots versus humans:

   (a) *Are bots more politically diverse than humans?* Using measures for the diversity of political attention, we consider whether bots consume information more frequently across the partisan aisle.
   (b) *Do bots and humans focus on the same issues and misinformation narratives?* We study what proportion of each cohorts discourse relate to certain topics.
   (c) *Do bots and humans frame critical COVID-19 issues and misinformation the same way?* Using semantic analysis, we contrast how bots and humans frame issues and if they are more susceptible to specific uses in language.

Method

Data set and user augmentation

We utilized the largest public COVID-19 Twitter data set, subset between January 21, 2020 and April 01, 2021 [17]. With more than 2 billion tweets collected, this allows us to analyze more than a years worth of online discourse regarding the pandemic. Given our interest in the USA and the predominantly anglo-based collection strategy, we restrict our analysis to only tweets in English.

As specified by our research questions, we are predominately interested in two aspects of each user. First, their political orientation and second their level of automation. Since we are interested in how the identity of users, rather than solely discourse, our data are subset on users whom we can generate data. We summarize the descriptive statistics of our dataset in Table 1. It should be noted as our labeling pipeline rely on external APIs, and hence subject to rate limits, we first sorted users based on their average number of tweets and retweets, to better capture the core and elite users within our analysis.

Automation detection

Bot detection is not only a fundamental method in assessing social media manipulation, but one that co-evolves with increasingly sophisticated bots. Identifying bots is

a methodological arms race, and as such, we employ the latest version of Botometer (version 4) [14]. Botometer is a machine learning-based application programming interface produced and maintained by Indiana University. Both the prior version and current version use an ensemble classifier (a probabilistic mixture of different machine learning classifiers) to output a botscore, which indicates the likelihood a user is a bot or a human. For each user, Botometer extracts over 1000 features per user using their most recent 200 tweets and the user's profile to make an inference. This includes an account's profile, friends, temporal activity, language, and sentiment. The models are trained then on human annotated users that indicate the likelihood a user is a social bot. For this paper, we use the raw scores—which indicates exactly this likelihood.

The distribution of Botscores within our dataset is visualized in Fig. 1, which shows a natural bimodal break between 0.5. Hence, we label users with botscores above 0.5 as bots.

However, since botometer was trained on a different dataset than the Covid-19 dataset, we conduct further manual validation to ensure the scores effectively transfer to the new domain. We randomly sampled 200 accounts (100 from each account) which were validated by three researchers. This yielded an average of 88% agreement with the Botometer on average. We are conducting further validation by increasing the number of accounts (500) and manual inference, which will end in March (in-time to be updated for the paper-ready revision).

Political affiliation inference

To categorize a user's information diet, we cross-reference the URLs shared by users with the MediaBias/Fact-Check database. Media-bias/Fact-check is an independent

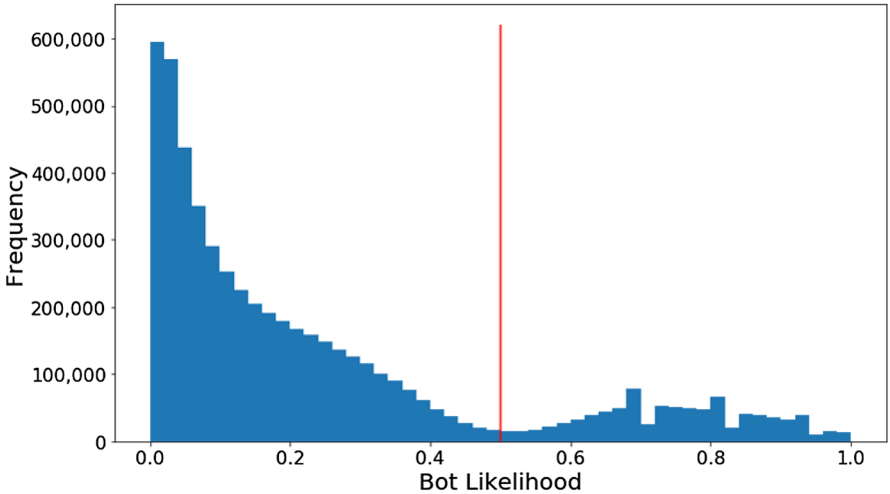

Fig. 1 Figure 1 shows the botscore distribution (which is the likelihood that a user is a bot), which forms a natural bimodal distribution break at 0.5

Table 1 Descriptive statistics of our data set

Statistic	Original amount in dataset	Subset in this paper
# of Users	42,124,899	5,499,308
# of Retweets	499,008,195	207,260,607
# Replies	80,249,435	23,327,314

fact-checking database which also rates the political leanings of media domains and websites. These are rated on a scale of five labels—left, center-left, center, center-right and right.

This process yielded a total of 97,810,067 posts with political URLs, and as a resultant 1,122,167 political users who have at least 5 URLs. We then used the label propogation algortihm to boost this to a total of 3,806,414 total users, which is larger than Luceri et al. by a factor of 100. We discuss the label propogation algorithm in detail in "Network analysis". The resultant data subset compared to the original dataset is thus given in Table 1.

We also take the time to justify a few of these choices—first, these retweets preclude quoted retweets, which may be used to counter opposing opinions. Metaxas and colleagues have found most retweets are endorsements, unless they arise from journalists or journalistic outlets. As the large majority of users are not journalists, results demonstrating endorsement, and thus affiliate political diet, would hold.

Similar to the validation of humans and bots, we randomly sample 100 users from each of these categories, with an average accuracy of 76% across three validators, through considering user descriptions and general tweeting behavior. Note, this is due to some accounts not having immediate identifiable hashtags or group-affiliations.In particular, Rathje and colleagues among other researchers have noted the growing tribalism online—identificaiton by particular group-tags (i.e. MAGA, TheResistance) [18]. As with the human-bot evaluation, we are increasing our validation to 300 per group across 5 validators.

Network analysis

We construct a network using our 5 million users. A network is defined as $G = (V, E)$, where V is a set of vertices and E a set of tuples containing $u, v \in V$ (i.e., $e = (u, v)$, $e \in E$. In our case, we model users as nodes and retweets as edges. That is, a directed edge is drawn from source A to retweeter B for every retweet. The weight of the edge then denotes the frequency. We also construct a second network based on replies. This allows us to contrast behavioral difference between retweets and replies, which will also be used in "Human–bot interaction".

Given the large-size and many users who tweet sporadically in low volumes, for visualization, we first conduct a K-core decomposition [19]. A k-core decomposition of graph G essentially iterative prunes nodes until every node has at least degree k. This allows analysis of the core, often elite users, within a network. Then, we use a force-based algorithm through Netwulf [20], a package built on the d3-force

module, which uses Verlet integration to calculate optimal positions for the n-body problem [21].

Further augmentation: measuring political extremism

For certain visualizations, we take their most frequent label to be their political affiliation. However, to investigate polarization and the diversity of attention, we also construct two measures for our research question. The first measures the political position of a user, the second measures the diversity of attention for each user.

- *Weighted political position p* For each user, we take the weighted average of the media they consume. First, we assign the labels left, center-left, center, center-right, and right with the values $-2, -1, 0, 1$, and 2 respectively. Let $U(i)$ be the domains of URLs that user i retweeted, and denote $v(j) \in [-2, -1, 0, 1, 2]$ as the political position of URL j. Then, the weighted political position p is given in Eq. 1 as:

$$p(i) = \frac{1}{|U(i)|} \sum_{j \in U(i)} v(j) \tag{1}$$

- *Diversity of attention* For every user, we consider the mean position of the users they retweet, which is equivalent to its neighbors based on incoming edges. We take the average of this value, take the absolute value of this value and the user's own political position, then subtract them. This gives the difference between the user and what they consume. Explicitly in Eq. 2:

$$\Delta(i) = p(i) - \frac{1}{|\mathrm{Nei}(i)|} \sum_{j \in \mathrm{Nei}(i)} p(j) \tag{2}$$

Here, $\mathrm{Nei}(i)$ denotes the neighbors of user i. Combined with centrality, the political position and dispersion of each user effectively tells us how embedded and polarized a user is. For liberals, if $\Delta(i) > 0$ is positive then a user is more likely to consume extreme material. Vice versa for conservatives, $\Delta(i) > 0$ implies each user on average is more conservative than their friends.

Human–bot interaction

To measure the effectiveness of bots engaging with humans, we borrow three measures from Luceri et al. [16]. They aim to measure the extent humans rely on and interact with content generated or facilitated by social bots. These are given by:

- *Retweet Pervasiveness (RTP)* measures the intrusiveness of bot-generated content:

$$\text{RTP} = \frac{\text{\# of human retweets from bot tweets}}{\text{\# of human retweets}} \tag{3}$$

- *Reply rate (RR)* measures the percentage of replies given by humans are to social bots.

$$\text{RR} = \frac{\text{\# of human replies from bot tweets}}{\text{\# of human replies}} \tag{4}$$

- *Human to Bot Rate (H2BR)* quantifies human interaction with bots over all human activity in the network.

$$\text{H2BR} = \frac{\text{\# of human interaction with bots}}{\text{\# of human interaction}} \tag{5}$$

We further modify these measures by dividing by the percentage of bots per group, as this gives a better sense of how bots perform against each other, pound-for-pound.

Discourse analysis

Drawing upon prior work that utilizes the same dataset, we focus on key events, themes, and pieces of misinformation identified within these studies [8, 22, 23]. We focus on four primary discourse topics and three pieces of misinformation. We provide a brief summary for each and the boolean search query and associated tokens we used to extract them. Since these topics and pieces of misinformation do not have significant overlaps with other topics, direct querying producing clean results. We validate the consistency of each of these topics by sampling 100 tweets each, with 96% average accuracy.

- *Discourse Topics:*

    1. *Masking* As mask wearing was not commonly accepted public health practices, the practice produced significant amounts of push back [24]. Query: "mask"
    2. *Vaccines* The deployment and development of vaccines was the center of attention and debate for much of the pandemic. Beginning with then President Trump's frequent declarations that vaccines (through Operation Warp Speed [24]) would be coming and the eventual roll out during President Biden's administration [25]. Query: "vaccin" or "vax"
    3. *Relief bill* During the pandemic, there were multiple relief bills that were considered by the federal government, starting with the *Coronavirus Preparedness and Repsonse SUpplmenetal APpropriations Act, 2020* [26]. Each of these, costing from $8.8 billion to $2.2 trillion were significant sources of debate [27]. Query: "relief" and ("bill or "covid").
    4. *Social distancing* Social distancing is a concept in public health where individuals keep away from each other to minimize transmission. In the early

stages when comparisons between COVID-19 and the flu were drawn, there was a lot of resistance toward the practice, especially with the help of then President Trump [28]. Query: “soc” and “distanc”

- *Misinformation* [9, 29]:

  1. *Bill Gates* Online trolls and misinformation communities have claimed the billionaire planned the pandemic so he could use the vaccine as an opportunity to inject trackers into everyone’s body. Query: “bill” and “gates”
  2. *Bio-weapon* Although scientists had yet to figure out the origins of the virus, there were online speculation that COVID-19 was a biological weapon intended for medical warfare and control over-population. Query: “bio” and “weapon”
  3. *5G* Conspiracy theorists believed that 5G towers were the conduit for COVID-19 virality, and that radio waves emitted by 5G towers cause mutations in the body and made us more vulnerable. Importantly, 5G celltower related misinformation has been around for more than a decade, with celebrity endorsements. Query: “5g”

Upon extracting all tweets related to these issues, we then process these tweets into tokens, removing extraneous tokens such as "rt" and stop words, then stem and lemmatize each word. This yielded more than 131 million unique tokens and thus, we focused only on the 500,000 most frequent words. As we will do for the prior analyses, we breakdown our independent variable as the product space of [humans, bots] and [liberals (left), conservatives(right)], then consider the top 100 words for each cohort. However, since many of these tokens will also overlap, for each topic we remove the intersection of these word lists—this allows us to analyze how each group frames the issue, based on the unique words they use.

Results

Bot political leanings

Since we have tagged users by their political affiliation and bot identity, we subset our cohorts into liberal humans, liberal bots, conservative humans, and conservative

Table 2 Breakdown of users in each group, across (a) all tweets and (b) all users

	Liberal	Conservative
	(a) Tweet-based statistics	
Humans	106,300,970 (58.3%)	33,664,302 (18.5%)
Bots	31,859,563 (17.4%)	10,534,638 (5.78%)
	(b) User-based statistics	
Humans	2,697,725 (72.9%)	587,679 (15.9%)
Bots	322,553 (8.7%)	93,621 (2.53%)

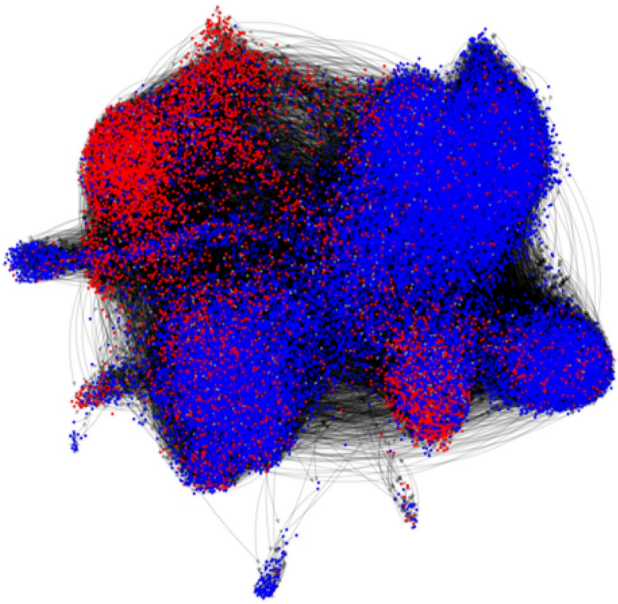

Fig. 2 20-Core graph of the COVID-19 Twitter Network

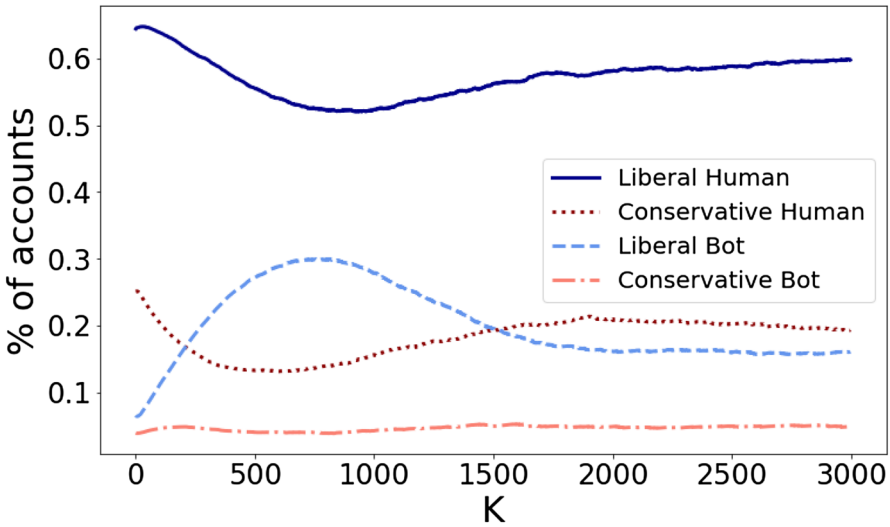

Fig. 3 The proportion of users iteratively over K-Core decompositions

bots. Immediately, we notice liberal humans greatly outnumber conservative humans. This is consistent with a 2019 Pew Center study that showed a ratio of 63% to 37% ratio between Democrats and Republicans (based on 36% vs 21% in raw percentages) Table 2.

More interesting is when comparing between tweet-based and user-based statistics. For instance, bots saturate almost twice as much at the tweet level compared to their overall population. Conservatives are also "louder" in that their user representation is slightly less than their tweet level representation. Similar to Luceri et al., we discover an imbalance in liberal (8.7% of users) versus conservative bots (2.53% of users) , but not as pronounced as liberal (72.9%) versus conservative (15.9%) humans. Figure 2 visualizes the network of users. We observe significant polarization across liberal and conservative users.

Figure 3 further shows how iterative K-core decompositions shift the distribution of users. Note, we use an unweighted graph here, which removes the effects of over-frequent tweeting by bots and focuses more on the network topology. We find liberal humans dominate the composition of COVID-19-related tweets. Conservative bots maintain a low presence. Interestingly, the presence of liberal bots increases between the 0 and 1000 core region and bound conservative humans from above until $k = 1500$. This may be due to the following reason. Bots that tweet about the news and diversely may occupy a bigger presence earlier on. However, as the core network becomes smaller and more elite, humans with external presence—such as health officials and senators—become more dominant.

We also conclude the effect of conservative bots are more muted when compared to election events [16]. We posit this may be due to two reasons. First, the weakness of conservative bots may be due to the pandemic still primarily being health-related issue rather than a political one, and such does not activate the same base. Second and related, the dataset tracks pandemic-related discourse explicitly, rather than political discourse, and thus may skew another way.

Lastly, we also consider how embedded humans and bots are within the network. Table 3 shows the centrality measures averaged for all users. In general, for both the centrality measures of the average bot is higher than the average human, likely due to many sporadic users. More interesting is the comparison between conservative and liberal users. We observe that conservative humans have higher centrality that liberal humans, which suggests a denser network structure and tighter community. In contrast, liberal bots are more embedded in the Twitter network than conservative bots.

Table 3 (a) Out-degree centrality and (b) in-degree centrality of the entire Twitter network

	Liberal	Conservative
(a) Out-degree centrality		
Humans	7.81×10^{-5}	8.82×10^{-5}
Bots	3.69×10^{-4}	2.56×10^{-4}
Ratio (B/H)	4.72	2.9
(b) In-degree centrality		
Humans	4.00×10^{-4}	5.10×10^{-4}
Bots	1.23×10^{-3}	1.09×10^{-3}
Ratio (B/H)	3.08	2.14

Table 4 Percentage human engagement over liberal and conservative bot tweets

	Liberal absolute	Conservative absolute	Liberal relative	Conservative relative
RTP	0.203	0.045	2.333	1.779
RR	0.108	0.044	1.241	1.739
H2BR	0.193	0.045	2.218	1.779

Absolute stands for total human-bot interaction over all human interaction. Relative is normalized over the number of bots

Human-engagement and bot effectiveness

Table 4 shows the percentage of human engagement for liberal and conservative bots. It shows the amount of human engagement in absolute terms (over the total number of human activity) and also relative terms (normalized over the percentage of bot accounts). In absolute terms, liberal bots dominate the amount of retweets, taking up to 20% of all retweets. In contrast, conservative bots only amount to around 4.5% of the total human retweets.

Interestingly, when comparing replies to humans, conservative bots perform much better. This implies red bots are better at eliciting repsonses from humans. This is significant as, if we assume humans reply to other "perceived" humans more frequently, then conservative bots are better at being perceived as humans.

Political analysis and human comparisons

Upon identifying the strategies found in bots, we compare the political position and influence of each of these groups. Table 5 show the average political position by cohort, the average political position of their neighbors, the difference in their position, and the standard deviation in the position. The standard deviation here is meant to capture the diversity in each user's neighbors.

We observe that interestingly, a user's average neighbor position regresses toward the center. This indicates that during the pandemic, each user sees on average a more centrist position than their own on Twitter. A corollary of this is that conservative bots and conservative humans appear to retweet more frequently across the aisle.

Table 5 Political positions of liberal and conservative humans and bots

	User pol	Neighbor ave. pol	Δ pol	Std Δ pol
Liberal bots	− 0.492	− 0.461	− 0.007	0.415
Liberal humans	− 0.500	− 0.462	− 0.010	0.424
Conserv. bots	0.549	0.315	0.160	0.595
Conserv. humans	0.469	0.261	0.162	0.563

The table summarizes the average political position by group, the average political position of their neighbors, the difference in position, and the standard deviation of the difference in position

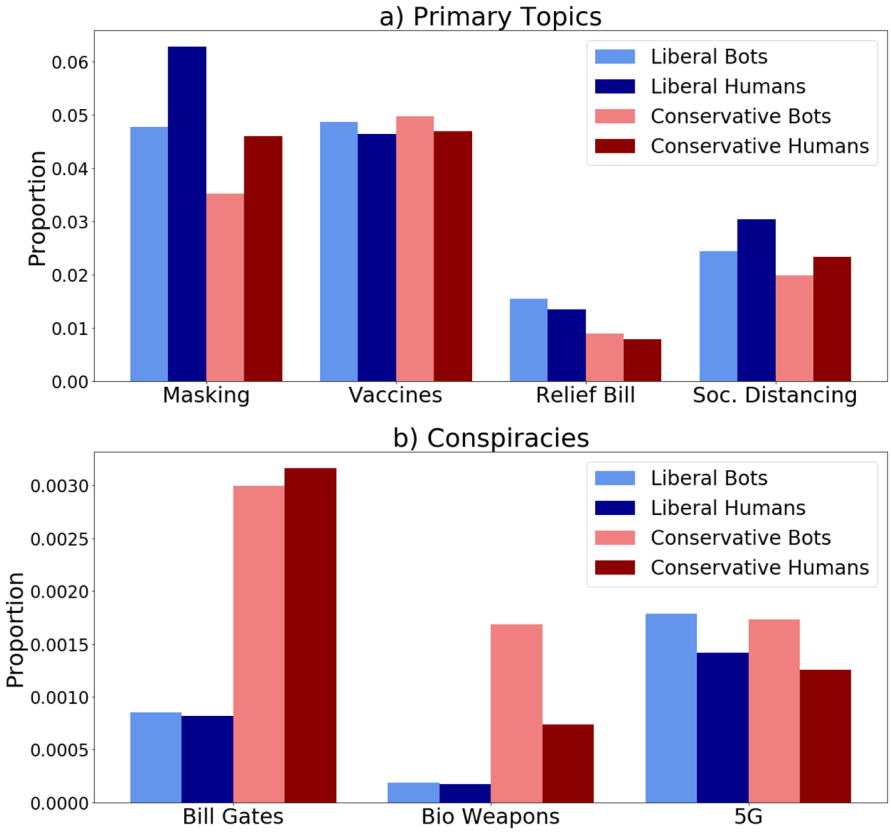

Fig. 4 Topic coverage by each group. **a** Covers primary topics—masking, vaccinations, relief bills, and social distancing. **b** Covers three significant misinformation narratives—COVID-19 was started by Bill Gates, COVID-19 is a bio-weapon, and COVID-19 is spread by 5G

This may be, however, a product of the left-leaning media's presence in pandemic coverage, as these are in opposition with prior studies strictly related to politics [16].

Topical analysis

Figure 4 shows the relative proportion for which each group engaged with top topics (masking, vaccines, relief bills, and social distancing) and misinformation (Bill Gates profiting from COVID-19, COVID-19 is a bio-weapon, and COVID-19 was spread through 5G. The proportions for mask wearing and social distancing near each other. Liberal humans take the lead at 6%, and in general humans discuss masquering and social distancing more than their bot counterparts. Vaccines in contrast are discussed in relatively equal proportions, and for the relief bill, bots tend to tweet more than humans. This suggests that bots are more in tune with political content, rather than health related content.

Table 6 Comparison of how core issues were framed during the COVID-19 pandemic

Topic	Position	Bot	Top words	Excluded top
Masks	Liberal	Bot	Mask, wear, distance, **trump**, social	**Wearamask**, save, republican, presid, joebiden
	Liberal	Human	Mask, wear, distance, social, **trump**	**Wearamask**, fuck, die, joebiden, save
	Conservative	Bot	Mask, wear, distanc, social, face	Fauci, force, **wearamask**, govern, busi
	Conservative	Human	Mask, wear, distance, social, **lockdown**	Force, studi, busi, fauci, scienc
Relief_Bill	Liberal	Bot	Relief, senate, **republican**, vote, pass	Mcconnel, mitch, rbreich, **right**, kylegriffin
	Liberal	Human	Relief, senate, **republican**, vote, pass	Mcconnel, mitch, rbreich, **right**, today
	Conservative	Bot	Relief, **democrat**, pelosi, biden, senate	Pelosi, realdonaldtrump, call, nanci, news
	Conservative	Human	Relief, **democrat**, pelosi, senate, vote	Pelosi, call, realdonaldtrump, nanci, send
Social distancing	Liberal	Bot	Distance, social, stay, home, mask	Stayathome, love, staysafe, friend, quarantine
	Liberal	Human	Stay, home, social, distance, mask	**Fuck**, love, sick, limit, health
	Conservative	Bot	Social, distance, stay, home, mask	Talk, **democrat**, biden, song, fantasticradio
	Conservative	Human	Social, distance, stay, home, mask	Close, biden, told, **democrat**, allow
Vaccines	Liberal	Bot	Vaccine, dose, pfizer, **trump**, health	Administer, free, china, nation, protect
	Liberal	Human	Vaccine, dose, pfizer, effect, **trump**	Free, good, protect, administer, hope
	Conservative	Bot	Vaccine, india, dose, pfizer, effect	China, gate, narendromodi, lockdown, fauci
	Conservative	Human	Vaccine, dose, pfizer, effect, need	Lockdown, passport, die, gate, china

The bolded text show comparisons we make per topic, given a specific topic

Indeed, when we discuss misinformation we see a vast increase in the activity of conservative bots and humans. Conservative users retweet misinformation about Bill Gates more than three times than their liberal counter parts. A similar increase is observed with the narrative that COVID-19 is a bio-weapon. Interestingly, the activity by conservative bots is more than twice their conservative human counterparts. In contrast, misinformation about 5G is relatively more ubiquitous across all user groups. This may be related to the existing history of 5G skeptics, who believe that radio waves can induce cancer and other diseases.

A natural follow-up question is why these differences in volume exist. In other words, how are these issues framed differently by these different groups. Table 6 compares how key topics were framed, showing the top five words for each cohort, and top five words that did not intersect with other groups. When discussing mask wearing, liberal and conservatives both mention similar topics. However both liberal bots and humans prioritize discourse about Trump. In contrast, conservative humans focus more on lockdown policies.

Differences clarify when we consider the top keywords in absence of the common words. Both liberal bots and humans mention Joe Biden. In contrast, conservative parts and humans focus on Fauci and being forced to comply with mask wearing and lock down policies. Along with business related discourse. However, perhaps the most interesting is how the #WearAMask Appears in all groups except for conservative humans. This suggests a deliberate avoidance of the term. When comparing liberal and conservative humans, we also noticed that liberal humans employ much more in civil language. This framing will prove important when considering the in-group and out-group dynamics in online commmunication, and the selective attention of bots depending on their political diet.

In regards to relief bill, the focus on our group negativity comes to life. Among the top terms for liberals, mentioning Republicans in the Senate is popular. Similarly, popular among conservatives is mentioning Democrats, Nancy Pelosi, and also Biden. This dynamic increases when we look at top terms excluding common terms, or Democrats reference Mitch McConnell whereas Republicans mention Nancy Pelosi. However they do also mention prominent in party members, such as Robert Reich for liberals, and Donald Trump for conservatives. For social distancing, the top words are in great alignment. Whereas liberals tend to focus on the stay at home discourse, conservatives discuss the opposing party more frequently. We also see a similar use of incivil language amongst human liberals, that are absent from all three other groups.

Moving forward, we aim to conduct a deeper analysis of framing and semantic content, through techniques such as SAGE or Log-likelihood lingusitic analysis specific to Covid-19 [30]. Moreoever, further inquiry into the identity of these bots—such as news accounts or ones playing "social good" roles would be of interesting further study.

Discussion

In this paper, we investigated how humans and bots interacted politically and discursively on Twitter during the COVID-19 pandemic. Using 5.4 million users as our sample size, we discovered a few key results. First, while liberal users dominate the discourse about COVID-19, bots tweet twice as much as their user-level proportions.

We find bots to be more central. However, this is likely due to the large amount of sporadic users and indeed, when considering iterative K-Core decompositions, the role of bots diminish. An interesting dynamic here is between the 0 and 1500, the proportion of liberal bots is higher than humans. We posit this may be due to the critical role bots have in sharing information, which is more significant at lower K-core subgraphs. At higher values of K, elite human members take over.

This is further evidenced by our analysis of human engagement. Whereas liberal bots elicit greater response in retweets—up to 20% of all human retweets—conservative bots are much better at eliciting replies. One hypothesis is conservative bots may be more easily "perceived" as humans. Certainly, the role of verified accounts, authenticity, and information flow certainly merits further investigation, such as through an experiment.

Interestingly, we find conservative users seem to retweet with more diversity, which diverges from what prior studies have shown regarding echo chambers during political events [31]. It is important to note that this is likely due to the *absence* of conservative bots. This is not to say their involvement on social media is less prominent, simply that their involvement in strictly pandemic related discourse is limited. Additionally, as we mentioned in "Method", our measure is more accurately defined as the political diet of users, which may be influenced by how the institutional media publishes about the pandemic.

Lastly, the most exciting results this paper offers are about how liberal and conservative humans and bots *framed* the pandemic differently. Humans pay more attention to masking and social distancing discourse, whereas bots engage with more political topics—such as the relief bills. Partisan asymmetries emerge in misinformation narratives. Conservative users tweet frequently about the Bill Gates and bio-weapon conspiracy, and conservative bots in particular promogulated the bio-weapon narrative up to 200% than their human counterparts. In contrast, both liberals and conservatives engaged with the 5G conspiracy.

In terms of semantics, conservative humans actively avoid the hashtag #wearAMask, and focuses more on lockdown policy discourse. This is critical as it shows evidentially how politics can modulate health practices. For the relief bills, all users have the propensity to tweet about the out-party (i.e., liberals tweet about conservatives, conservatives tweet about liberals). Social distancing appear as a health practice for liberals and a policy issue for conservatives. Lastly, liberal humans tend to use more incivil language especially when discussing health practices. This is in agreement with work by Druckman et al. which show conservatives tend to avoid incivil language [32].

The greatest importance of these findings has to do with the flow of information. Given the centrality measures and our pervasiveness metrics, this would suggest

bots play some intermediary role between elites and everyday users. They do not set the agenda, but do much to distribute. Furthermore, there are asymmetries, such as the ability of conservative bots to engage with humans via replies. We also observe bots appear more sensitive to popular and controversial topics, and a follow-up study that ties popularity and topic sentiment with the sensitivity of bots would be of great interest.

In sum, we find a surprising absence of conservative bots in COVID-19 discourse compared to political conversations [33], and the selective absence of certain words as modulated by political orientation. Humans are better at selective framing. The primary limitation to this study is that bot-labeling is still ongoing. Although we have captured more than 50% of all retweets, this is from the most prolific tweeters and thus may skew results toward elite users. Additionally, refinement to how we treat political diet labeling can inform future research about the interaction of institutions, bots, and information flow.

Declarations

Conflict of interest On behalf of all authors, the corresponding author states that there is no conflict of interest.

References

1. Ferrara, E., Varol, O., Davis, C., Menczer, F., & Flammini, A. (2016). The rise of social bots. *Communications of the ACM, 59*(7), 96–104.
2. Boichak, O., Jackson, S., Hemsley, J., & Tanupabrungsun, S. (2018). Automated diffusion? bots and their influence during the 2016 us presidential election. In *International conference on information* (pp. 17–26). Springer.
3. Bovet, A., & Makse, H. A. (2019). Influence of fake news in twitter during the 2016 us presidential election. *Nature Communications, 10*(1), 1–14.
4. Badawy, A., Lerman, K., & Ferrara, E. (2019). Who falls for online political manipulation? In *Companion proceedings of the 2019 world wide web conference* (pp. 162–168).
5. Bail, C. A., Argyle, L. P., Brown, T. W., Bumpus, J. P., Chen, H., Hunzaker, M. F., Lee, J., Mann, M., Merhout, F., & Volfovsky, A. (2018). Exposure to opposing views on social media can increase political polarization. *Proceedings of the National Academy of Sciences, 115*(37), 9216–9221.
6. Flaxman, S., Goel, S., & Rao, J. M. (2016). Filter bubbles, echo chambers, and online news consumption. *Public Opinion Quarterly, 80*(S1), 298–320.
7. Druckman, J. N., Ognyanova, K., Baum, M. A., Lazer, D., Perlis, R. H., Volpe, J. D., et al. (2021). The role of race, religion, and partisanship in misperceptions about covid-19. *Group Processes and Intergroup Relations, 24*(4), 638–657.
8. Jiang, J., Chen, E., Yan, S., Lerman, K., & Ferrara, E. (2020). Political polarization drives online conversations about covid-19 in the united states. *Human Behavior and Emerging Technologies, 2*(3), 200–211.
9. Chen, E., Chang, H., Rao, A., Lerman, K., Cowan, G., & Ferrara, E. (2021). Covid-19 misinformation and the 2020 us presidential election. *The Harvard Kennedy School Misinformation Review*.
10. Chang, H.-C.H., Chen, E., Zhang, M., Muric, G., & Ferrara, E. (2021). Social bots and social media manipulation in 2020: The year in review. In U. Engel, A. Quan-Haase, X. Liu, & L. Lyberg (Eds.), *Handbook of computational social science* (p. 18). London: Routledge.
11. Barbara, V. (2021). *Miracle cures and magnetic people. Brazil's fake news is utterly bizarre*. New York Times.

12. González-Bailón, S., & De Domenico, M. (2021). Bots are less central than verified accounts during contentious political events. *Proceedings of the National Academy of Sciences, 118*(11).
13. Mønsted, B., Sapieżyński, P., Ferrara, E., & Lehmann, S. (2017). Evidence of complex contagion of information in social media: An experiment using twitter bots. *PloS One, 12*(9), 0184148.
14. Sayyadiharikandeh, M., Varol, O., Yang, K.-C., Flammini, A., & Menczer, F. (2020). Detection of novel social bots by ensembles of specialized classifiers. In: *Proceedings of the 29th ACM international conference on information & knowledge management* (pp. 2725–2732).
15. Druckman, J. N., Peterson, E., & Slothuus, R. (2013). How elite partisan polarization affects public opinion formation. *American Political Science Review, 107*(1), 57–79.
16. Luceri, L., Deb, A., Badawy, A., & Ferrara, E. (2019). Red bots do it better: Comparative analysis of social bot partisan behavior. In *Companion proceedings of the 2019 world wide web conference* (pp. 1007–1012)
17. Chen, E., Lerman, K., Ferrara, E., et al. (2020). Tracking social media discourse about the covid-19 pandemic: Development of a public coronavirus twitter data set. *JMIR Public Health and Surveillance, 6*(2), 19273.
18. Rathje, S., Van Bavel, J. J., & van der Linden, S. (2021). Out-group animosity drives engagement on social media. *Proceedings of the National Academy of Sciences*, *118*(26).
19. Alvarez-Hamelin, J.I., Dall'Asta, L., Barrat, A., & Vespignani, A. (2006). Large scale networks fingerprinting and visualization using the k-core decomposition. In *Advances in neural information processing systems* (pp. 41–50).
20. Aslak, U., & Maier, B. F. (2019). Netwulf: Interactive visualization of networks in python. *Journal of Open Source Software, 4*(42), 1425.
21. Hairer, E., Lubich, C., & Wanner, G. (2003). Geometric numerical integration illustrated by the störmer-verlet method. *Acta Numerica, 12,* 399–450.
22. Brennen, J. S., Simon, F., Howard, P. N., & Nielsen, R. K. (2020). Types, sources, and claims of covid-19 misinformation. *Reuters Institute, 7*(3), 1.
23. Roozenbeek, J., Schneider, C. R., Dryhurst, S., Kerr, J., Freeman, A. L., Recchia, G., et al. (2020). Susceptibility to misinformation about Covid-19 around the world. *Royal Society Open Science, 7*(10), 201199.
24. Javid, B., Weekes, M. P., & Matheson, N. J. (2020). Covid-19: Should the public wear face masks? *British Medical Journal Publishing Group*.
25. LaFraniere, S. (2021). *Biden got the vaccine rollout humming, with trump's help*. New York Times.
26. King, J. S. (2020). Covid-19 and the need for health care reform. *New England Journal of Medicine, 382*(26), 104.
27. Wu, N., & Zarracina, J. (2021) All of the COVID-19 stimulus bills, visualized. Gannett Satellite Information Network. https://www.usatoday.com/in-depth/news/2021/03/11/covid-19-stimulus-how-much-do-coronavirus-relief-bills-cost/4602942001/.
28. Shear, M.D., & Mervosh, S. (2020). *Trump encourages protest against governors who have imposed virus restrictions*. New York Times, April 18.
29. Weiss, S., & Greenstreet, S. (2020). *Coronavirus conspiracy theories don't stop at Bill Gates and 5G*. New York Post. https://nypost.com/2020/04/24/the-top-5-coronavirus-conspiracy-theories-bill-gates-5g-more/.
30. Smith, A., Tofu, D. A., Jalal, M., Halim, E. E., Sun, Y., Akavoor, V., Betke, M., Ishwar, P., Guo, L., & Wijaya, D. (2020). *Openframing: We brought the ml; you bring the data. interact with your data and discover its frames*. arXiv preprint arXiv:2008.06974.
31. Wojcieszak, M., Casas, A., Yu, X., Nagler, J., & Tucker, J. A. (2021). Echo chambers revisited: The (overwhelming) sharing of in-group politicians, pundits and media on twitter.
32. Druckman, J. N., Gubitz, S., Lloyd, A. M., & Levendusky, M. S. (2019). How incivility on partisan media (de) polarizes the electorate. *The Journal of Politics, 81*(1), 291–295.
33. Ferrara, E., Chang, H., Chen, E., Muric, G., & Patel, J. (2020). Characterizing social media manipulation in the 2020 us presidential election. First Monday.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.