

RESEARCH ARTICLE

Influence of augmented humans in online interactions during voting events

Massimo Stella[‡], Marco Cristoforetti, Manlio De Domenico^{‡*}

Fondazione Bruno Kessler, Via Sommarive 18, 38123 Povo (TN), Italy

[‡] Current address: Complex Science Consulting, Via A. Foscarini 2, 73199 Lecce (LE), Italy* mdedomenico@fbk.eu

OPEN ACCESS

Citation: Stella M, Cristoforetti M, De Domenico M (2019) Influence of augmented humans in online interactions during voting events. *PLoS ONE* 14(5): e0214210. <https://doi.org/10.1371/journal.pone.0214210>

Editor: Katarzyna Musial-Gabrys, University of Technology Sydney, AUSTRALIA

Received: July 23, 2018

Accepted: March 7, 2019

Published: May 16, 2019

Copyright: © 2019 Stella et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: Publicly available tweets were used for this study and the details of the search parameters used to collect the tweets through the Twitter API are described in the Methods section. Due to recent Twitter updates to the policies (Developer Agreement and Policy: <https://developer.twitter.com/en/developer-terms/more-on-restricted-use-cases.html>), the IDs of the tweets collected in this study are not shared on a public repository but are available upon request to the corresponding author or co-author, Marco Cristoforetti at mcristofo@fbk.eu. Sensitive information such as political affiliation or beliefs will

Abstract

The advent of the digital era provided a fertile ground for the development of virtual societies, complex systems influencing real-world dynamics. Understanding online human behavior and its relevance beyond the digital boundaries is still an open challenge. Here we show that online social interactions during a massive voting event can be used to build an accurate map of real-world political parties and electoral ranks for Italian elections in 2018. We provide evidence that information flow and collective attention are often driven by a special class of highly influential users, that we name “augmented humans”, who exploit thousands of automated agents, also known as bots, for enhancing their online influence. We show that augmented humans generate deep information cascades, to the same extent of news media and other broadcasters, while they uniformly infiltrate across the full range of identified groups. Digital augmentation represents the cyber-physical counterpart of the human desire to acquire power within social systems.

Introduction

Online social actions drive collective attention and dynamics [1, 2], having a deep impact on the construction and perception of social reality. Many large-scale studies have reported evidence of online ecosystems altering decision-making of crowds [3] and influencing real-world voting of millions of people [4]. The last few years have seen a deluge of increasingly more sophisticated automated online agents, called also “bots”, populating techno-social systems cleverly disguised as human users [5, 6, 7, 8, 9]. Nowadays, bots can produce credible content with human-like temporal patterns [10, 11, 12]. By promoting online activity, bots can interact with humans and influence their standing against specific topics such as political issues [10, 7, 12, 13]. Since manoeuvring social platforms can deeply affect real-world dynamics [14, 15], understanding if and how computer-generated activities can alter the behavioral responses of humans to achieve online social manipulation is of utmost importance [16, 5, 17]. Identifying and quantifying these effects is particularly crucial during voting events, where individuals’ decisions might be driven by external events, such as natural disasters or economic shocks [18]. While attention is generally paid to how physical interactions among voters and electoral arrangements influence voting behavior, Bruter and Harrison [19] shifted the focus on the

not be distributed with the tweets IDs, to comply with Twitter policies on data protection and redistribution. Requesters will be asked to sign a declaration of purposes for academic/non-commercial use of the data and to comply with Twitter's Developer Agreement and Policy on data protection and redistribution.

Funding: The author(s) received no specific funding for this work.

Competing interests: The authors have declared that no competing interests exist.

psychological influence that electoral arrangements exert on voters by altering human emotions and behavior. The investigation of voting from a cognitive perspective leads to the concept of *electoral ergonomics*: Understanding the optimal ways in which voters emotionally cope with voting outcomes can lead to a better prediction of the elections.

Here we quantify to which extent online social activity reflects the real world by considering a data driven approach using streaming data from social media for analysing microscopic patterns between users, an increasingly common approach in the computational social sciences [20, 21]. We characterize the peculiar behavior of a class of individuals who make a massive use of bots to enhance their online visibility and influence. The term *cyborg* has been used in this context to identify, indistinctly, bot-assisted human or human-assisted bot accounts generating spam content over social platforms such as Twitter [5, 22]. Here, we prefer to use the term *augmented human* for indicating specifically those human accounts exploiting bots for artificially increasing, i.e. *augmenting*, their influence in the digital world, analogously to physical augmentation improving human performances in the real world [23]. Like several automated agents identified in our data set, augmented humans played a special role for information spreading, by triggering deep information cascades with the help of bots.

Methods

Data collection

Between 24 February 2018 and 7 March 2018, we have collected 966,483 messages (tweets) posted by 194,273 different users to the microblogging platform Twitter, containing at least one of the following keywords or hashtags: “elezioni”, “#elezioni”, “#elezioni2018”, “#elezioni4-marzo”, “#ItalyElection2018”, “#voto”, “#4marzo”, “#M5S”, “#PD”, “#LeU”, “#LiberieUguali”, “#ForzaItalia”, “#FDI”, “#FI”, “#lega”, “#FratellidItalia”, “#MDP”.

Tweets have been collected using the streaming real-time provided by Twitter API platform, filtered by the above keywords. Twitter by default limits to 1% of the overall number of Tweets per second the fraction of tweets that can be retrieved from the streaming API. However, when the fraction of tweets concerning specific keywords is smaller than 1% of the global volume, Twitter does not apply limitations and the complete flow of information is collected. When this is not the case, Twitter provides messages of warning, reporting the cumulative number of missed tweets.

In the case of Italian elections, we received no warnings, therefore we have collected 100% of tweets containing the specified keywords.

We complied with the Twitter's terms of service to collect the data.

Classification task

In this work the classification of users in our data set as “humans” or “bots” is based on features providing the best classification accuracy according to recent studies [11, 12]: 1) *Statuses count*; 2) *Followers count*; 3) *Friends count*; 4) *Favourites count*; 5) *Listed count*; 6) *Default profile*; 7) *Geo enabled*; 8) *Profile use background image*; 9) *Protected*; 10) *Verified*. The total number of features is ten ($N_{feats} = 10$).

Searching for better performance we tested different machine learning techniques on an independent dataset created ad-hoc (see Supplementary Information) from a collection of manually annotated datasets (see Table 1). Models are trained on the 80% of the data and validated over the remaining 20%. The subdivision between the two sets was carried respecting the balancing between bots and humans present at the level of the single original datasets, in this way we have all type of different bots both in training and validation. The models based on

Table 1. Proportions of bot and human users in the training data.

data set	bot	human	total
cresci2015	0	5301	5301
cresci2017	7543	3474	11017
cyborgs	2756	0	2756
aboutme	0	2463	2463
omnibots	3530	0	3530
russian-trolls	389	0	389
TOTAL	14218	8775	22993

<https://doi.org/10.1371/journal.pone.0214210.t001>

random forest and deep neural network provided us with the highest accuracy (>90%) and precision in identifying bots (>95%). We chose the deep neural network model because it also provided a more stable classification of certain users playing the role of broadcasters (see Supplementary Information).

Fiedler partitioning, modularity, segregation and infiltration

Fiedler partitioning is a widely used technique from spectral graph theory for solving the min-max cut problem, i.e. partitioning a network in two components of similar size but connected by links whose total weights are the smallest possible [24]. Fiedler partitioning is obtained by considering the eigenvalue problem:

$$(D - W)\mathbf{q} = \lambda\mathbf{q}, \tag{1}$$

for a connected network represented by the weighted adjacency matrix W , with w_{ij} equal to the weight of the link between nodes i and j , and by a matrix D having the strength of nodes on its main diagonal. The spectral partitioning is obtained by identifying nodes relative to positive and negative entries in the second eigenvector \mathbf{q}_2 relative to the second eigenvalue λ_2 . q_2 and λ_2 are also called Fiedler vector and Fiedler value, respectively.

We use modularity [25] for identifying the polarization of users in the social bulk in two groups, labelled here by c_1 and c_2 :

$$\Phi_F = \frac{1}{2m} \sum_{ij} \left[A_{ij} - \frac{s_i s_j}{2s} \right] \delta_{c_i, c_j}. \tag{2}$$

Here, A_{ij} is 0 if users i and j did not interact, otherwise it is equal to the number of their interactions; s_i indicates the total number of interactions involving the i -th user, i.e. its strength, while s is the total number of interactions in the network. Polarization values Φ_F close to 0 indicate no antagonism between opposing factions, while Φ_F close to 1 is relative to strongly opposing factions.

We use the generalization of modularity to more than two groups for establishing the fragmentation of users in antagonizing social groups. The mathematical definition is similar to Eq (2), except for the fact that we consider more possible partitioning into a number communities (c_1, c_2, \dots, c_M) larger than 2. The number M of existing communities is not known *a priori* and an optimization process must be employed to discover best partitioning of the system.

We measure social segregation by considering the average size of connected components weighted by the number of their links. Indicating with \mathcal{C} the set of connected components and with \mathcal{C}_i the set of n_i nodes connected by e_i edges in the i -th connected component, we define

social segregation as:

$$\Sigma = 1 - \frac{\sum_{i=1}^{|\mathcal{C}|} n_i e_i}{\sum_{i=1}^{|\mathcal{C}|} n_i \times \sum_{i=1}^{|\mathcal{C}|} e_i} \tag{3}$$

Σ ranges between 0 (a network with a single connected component) and 1 (a network of isolated nodes with no links).

We define infiltration of a given type of users in a given social group i as the fraction of users of that type in group i , namely:

$$\mathcal{I}_s = \frac{\sum_{i=1}^{\mathcal{M}} s_i}{u_i} \tag{4}$$

where \mathcal{M} is the number of groups, s_i is the number of accounts of class s in the i -th group and u_i is the number of users in that group.

Results

Bot identification

To identify automated agents in the data set, we developed a deep neural network model (see [Methods](#) and [SI](#)), which classified 13.4% of users as bots, a value compatible to estimations during other voting events [[10](#), [11](#), [12](#)]. We built the network of interactions between human users and bots, including different types of social actions such as Retweets (i.e. a user sharing another user’s message), Mentions (i.e. a user mentioning another use in a message) and Replies (i.e. a user starting a discussion with another user). While Mentions and Replies can have both negative and positive connotations, Retweets are traditionally considered as a form of social endorsement [[26](#), [17](#)]: Users tend to retweet and thus endorse content they agree with.

Human-bot interactions: Homophily and centrality

[Fig 1](#) shows the volumes for messages (i.e., Tweets) and the considered social actions for both bots and humans. [Fig 1\(a\)–1\(d\)](#) indicates the overall fraction of messages exchanged between bots and humans (a), and the fractions stratified by social interactions (b-d). Arrows go from the source to the recipient of an interaction, for instance user A (source) replying to user B (recipient) would be indicated with an arrow $A \rightarrow B$. Most of the social interactions are from humans to bots (46%); Humans tend to interact with bots in 56% of mentions, 41% of replies and 43% of retweets. Bots interact with humans roughly in 4% of the interactions, independently on interaction type. This indicates that bots play a passive role in the network but are rather highly targeted by humans. [Fig 1\(e\)](#) shows the number of social interactions over time. The circadian rhythm is evident, i.e. at night the volume of messages generated by humans drops down. Also bots display a similar circadian rhythm, in agreement with previous observations [[10](#), [12](#)]. In general, bots contribute to 6% of the total number of social interactions occurred during the voting event (March 4 2018). [Fig 1\(f\)](#) reports the geographic locations of both human and bot users in the social system. Although most of the users are located in Italy, there are significant fractions of human users also located in the United States and in Europe, indicating the worldwide relevance of the Italian voting. Similarly, bots’ locations are

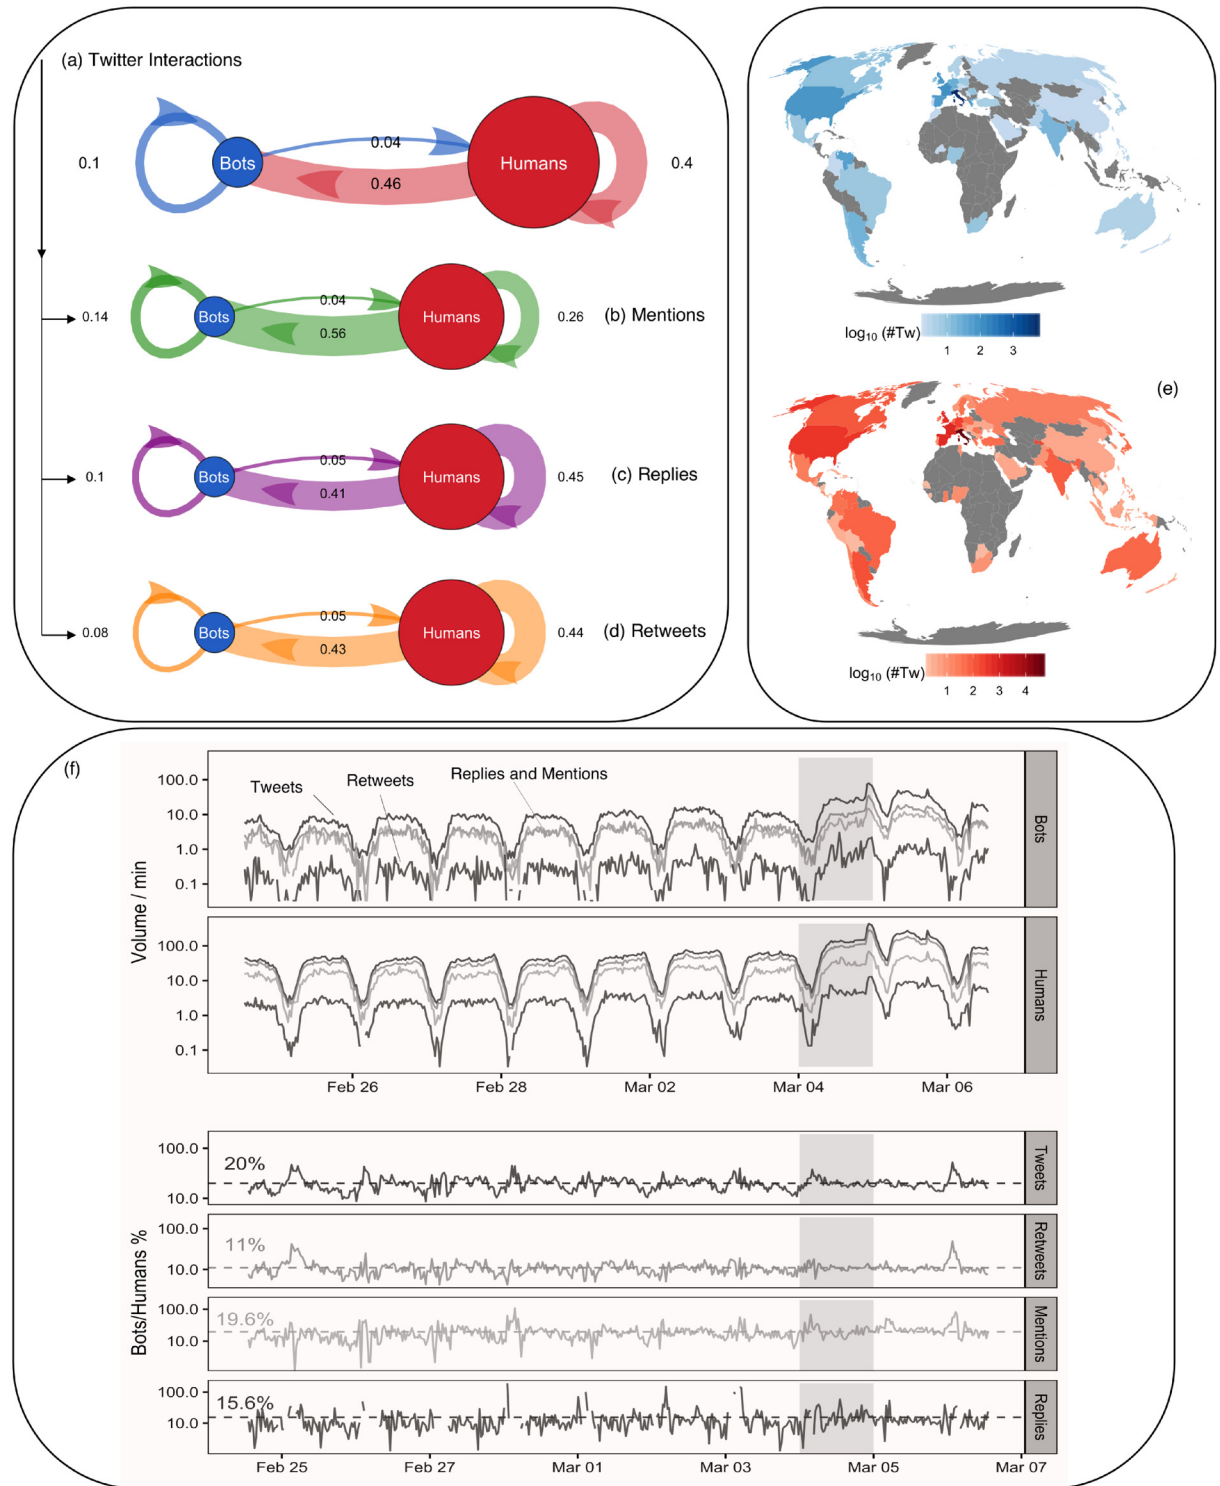


Fig 1. Online human-bot interactions during the Italian elections. (a): Volumes of human-bot interactions in Twitter. (b-d): Human-bot interactions stratified by actions: Mentions, Replies and Retweets. (e): Geographic location of involved users, where the color encodes the number of tweets per country, in logarithmic scale. As in (a), humans are in red and bots are in blue. Users are mostly located in Italy, with relevant interactions from other countries worldwide. (e): Evolution across time of the overall social activity of humans and bots (top), also stratified by actions (bottom).

<https://doi.org/10.1371/journal.pone.0214210.g001>

distributed worldwide and they are present in areas where no human users are geo-localized such as Morocco, Peru, Finland or Indonesia.

The analysis of observed social interactions (links) between users (nodes) before, during and after the voting day revealed bot homophily, i.e., automated agents tend to interact more with other bots rather than with humans compared to random expectation (see SI). Since interactions encode content spread [16], this result indicates that bots share messages mainly with each other and hence can resonate with the same content, be it news or spam. Furthermore, if we quantify the centrality of a user in terms of the probability of finding it by exploring the web of interactions at random, then we find that bots are almost twice as central as humans (see SI). The above findings indicate that bots play the role of sinks for information flow. In fact, 9 out of 10 hubs—i.e., highly interacting users—are bots and they are mainly news media and public profiles of politicians, which usually act as broadcasters and drive online information flow [17, 15]. The analysis of topic frequency and associations in bot-generated messages confirms this trend: Bots act as broadcasters by repeating the same political content of human users, boosting the spread of hashtags related to the electoral process (see SI).

Information cascades identify different classes of influencers

The observed social interactions build a complex network with a heterogeneous connectivity distribution. Such systems are well known for being susceptible to cascading events [27, 28] and, in the case of online social networks, the phenomenon might manifest as collective action and faster diffusion of specific information [29, 1, 2, 30]. In information cascades, a single piece of information is originated by a seed user, it is endorsed by other users in his/her neighborhood and consequently re-shared across the network [29]. Cascade size depends on a variety of factors [31], including—but not limited—to the structure of the network and the information content, making their prediction rather difficult [30].

Since metadata provided by Twitter does not allow to fully reconstruct an information cascade—because of missing intermediary retweets—we are only able to measure the cascade size by the overall number of endorsements, i.e. retweets, received by each post. Therefore, cascades are represented by star networks and, for brevity, in the following we will simply refer to them as “cascades”.

We have tracked 83,593 information cascades during Italian elections and, for each one, we have analyzed the underlying structure by measuring its size, i.e. the number of times an information has been re-shared. As expected for complex networks with highly heterogeneous connectivity [27], the distribution of observed cascade sizes is heavy-tailed and compatible with a power law characterized by a scaling exponent $\gamma = -2.33 \pm 0.04$, similarly to size distribution in percolation theory or avalanches in self-organized criticality [32]. Cascade size ranges between 2 and 4,313.

We show in Fig 2(a) a heat map of cascade size vs. the size of initiators' social neighborhood (i.e., the number of followers). As expected, on average, larger the number of followers larger the cascade size, with very few exceptions. Fig 2(b) shows the same data, with explicit information about user classification. This figure shows a good separation between human and bot behavior. Deeper information cascades are generated mostly by humans with a high number of followers, with the remarkable exception of one, User01, who produced the largest cascade among humans and bots despite having less than 100 followers.

Recently, dynamical activity-connectivity maps based on network and temporal activity patterns—or their variation—have been used to identify influential individuals or broadcasters during online protest diffusion [15] and contagion dynamics of extremist propaganda [33]. For instance, Bastos and Mercea [34] used hashtag trends for showing the existence of “serial

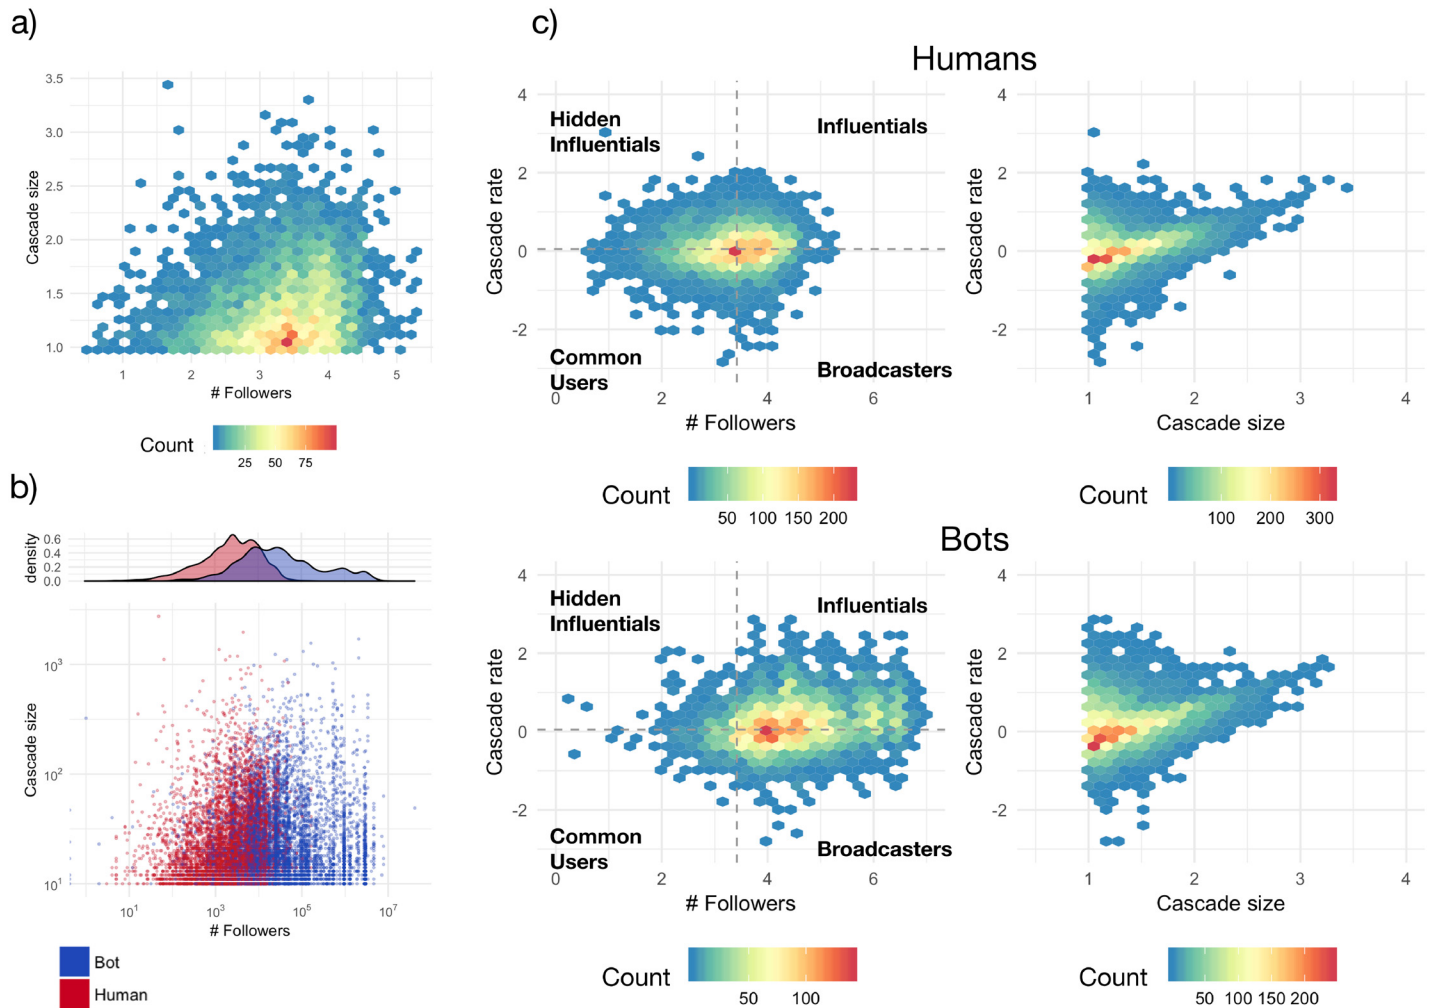


Fig 2. Information cascades during Italian elections. a) Heatmap of the number of users initiating information cascades, as a function of the size of their social neighborhood (Followers) and the size of the generated cascade; b) Scatter plot of the same data, with points encoding users. Color encodes bot/human classification and size encodes cascade’s diameter; c) As in a) but considering cascade rate (in units of retweets/hour), defined by the ratio between cascade size and its duration, vs. neighborhood size (left panels) and cascade size (right panels), for humans (top panels) and bots (bottom panels). The heatmap of cascade rate vs. neighborhood size allows one to identify 4 categories: hidden influentials, influentials, common users and broadcasters (see the text for further detail). Dashed lines indicate medians of structural and dynamical features in humans. Only cascades with at least 10 adopters are considered and, for heatmaps, the logarithm of the corresponding variables is considered.

<https://doi.org/10.1371/journal.pone.0214210.g002>

activists”, users with ordinary numbers of followers but very prolific in producing content about multiple political topics and bridging together disparate communities. Gonzalez et al. [15] related topological properties, such as the ratio between incoming (friends) and outgoing (followers) connections, to dynamical properties, such as the ratio of received and posted messages.

Here, we argue that it is also plausible to relate individuals’ social influence to the size of information cascades they generates with their content [35]. To this aim, we propose a more complex map relating a topological feature, i.e. the number of outgoing connections (followers), and a dynamical feature, i.e. information cascade growth rate, defined by the ratio between a cascade size and its duration over time. Baseline social behavior during a specific event, such as the Italian election in our case, is defined by the medians of the two observables, like shown in Fig 2(c). This map allows to easily identify four categories of individuals in the

social dynamics: i) *hidden influentials*, generating information cascades rapidly spreading from a small number of followers; ii) *influentials*, generating information cascades rapidly spreading from a large number of followers; iii) *broadcasters*, generating information cascades slowly spreading from a large number of followers; iv) *common users*, generating information cascades slowly spreading from a small number of followers. Remarkably, the topological and dynamical behavior of humans and bots is very different: during Italian elections, bots are mostly broadcasters (mostly media) and influentials (mostly political leaders). Fig 2(c) (right) highlights a positive correlation between cascade rate and size: Cascades involving more users tend also to flow over the interactions web at faster rates. This positive trend is stronger for cascades of sizes larger than 10^2 . The stronger correlation for larger cascades suggests that they differ qualitatively from smaller cascades: Larger cascades contain specific semantic content, in this case political-related topics, which accelerate spreading.

The social bulk of endorsements mirrors political antagonism

So far, our analysis characterized online human behavior in terms of human-bot interactions and information spreading. However, to quantify to which extent the observed online social activity reflects the real world a more sophisticated analysis is required. To this aim, we analyzed the static representation of the system, where interactions across time have been aggregated to a directed and weighted social network. We then identified the core of the observed social system by tracking the most relevant interactions among the most important actors. We identified relevant interactions by assuming that if two users share similar political ideologies, they can endorse and subsequently share (i.e. retweet) the content of each other. However, if only re-sharing was considered, the network would contain a lot of spurious connections due, for instance, to fortuitous endorsement rather than to a systematic intention.

We first filtered the network by considering only pair of users with at least one retweet, with either direction, because re-sharing content it is often a good proxy of social endorsement [26]. We then considered a more selective restriction, by requiring that at least another social action—i.e., either mention or reply—must be present in addition to a retweet. This restrictive selection allows one to filter out all spurious interactions among users with the advantage of not requiring any threshold with respect to the frequency of interactions themselves. The resulting network is what we call the *social bulk*, i.e. a network core of endorsement and exchange among users. By construction, information flows among users who share strong social relationships and are characterized by similar ideologies: In fact, when a retweet goes from one user to another one, both of them are endorsing the same content, thus making non-directionality a viable approach for representing the endorsement related to content sharing. Therefore, in the following, we can safely consider undirected interactions among users. Connections between users are weighted by the aggregated frequency of their social interactions. An illustration of how the social bulk is built is shown in Fig 3(a).

In the following, we introduce different measures to quantify different features of the social bulk, i.e. social polarization, fragmentation and segregation.

The concept of *social polarization* assumes the existence of two competing stances or opposing groups characterizing the mesoscale organization of the system [36]. In presence of two groups, they can be identified by calculating, for instance, the Fiedler partitioning [24, 12], which is related to the min-max cut problem for finding optimal flows in networks [24]. Fiedler partitioning (see Methods) separates the users of a connected graph into two classes such that the total number of inter-class connections is close to the optimal minimum. If interactions encode strong social relationships, as in the social bulk, then the Fiedler partitioning identifies two factions antagonizing each other by sharing the least endorsements possible.

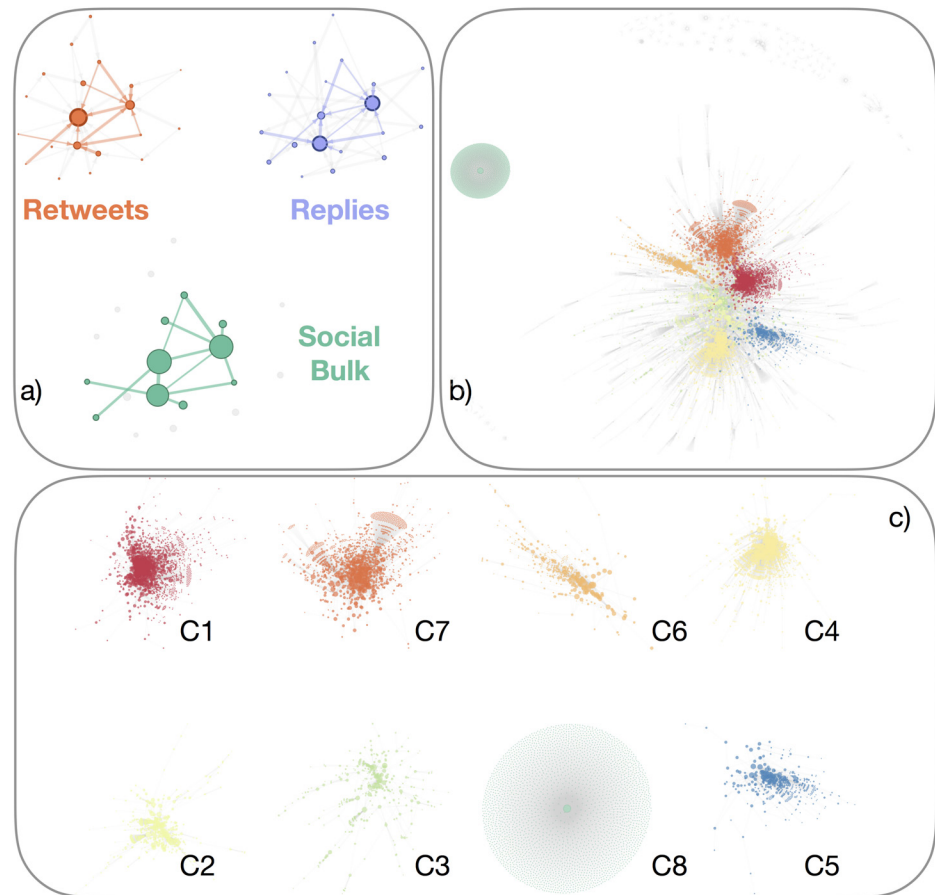


Fig 3. Social bulk of Italian elections. a) Twitter users can retweet or mention or reply with each other. Each action encode a specific social meaning and, by considering the co-existence of endorsement (i.e. retweet) and discussion (i.e. mention or reply), between the same pairs of users, we filter out spurious interactions to identify the social bulk of the system. b) Visualization of the social bulk emerged during Italian elections, with users (i.e., the nodes) colored by the community they belong to (see the text for further detail). c) The eight communities with at least 2% of users are represented separately, while preserving their relative position in the social bulk shown in panel b). Note the remarkable star-like topology C8 characterizing the augmented human identified in the system.

<https://doi.org/10.1371/journal.pone.0214210.g003>

We measure social polarization by computing the modularity [25] of the social bulk with respect to its Fiedler partitioning (see [Methods](#)). The larger the modularity Φ of the Fiedler partitioning, i.e. system polarization, the more antagonized are users into two opposing groups. For the largest connected component of the social bulk we calculate the polarization $\Phi_F = 0.452$. The expected polarization of a null model—where social relationships are uniformly randomized while preserving the individual degrees and the distribution of strengths—is $\langle \Phi_F^{rand} \rangle = 0.301$, significantly different from the observed network (p-value $< 10^{-5}$). This result indicates that the heterogeneity of social interactions can not explain, alone, the observed level of polarization, which has rather to be attributed to other causes such as political parties or opposing political trends. Notice that the social bulk analyzed here displays a higher modularity than other well studied social networks such as the Zachary Karate Club ($\Phi_F = 0.371$) and the dolphin’s social network ($\Phi_F = 0.401$)(cf. [37]). Considering that especially the Zachary Karate Club considers social interactions among two different groups, led by two different leaders and hence already highly polarized, then the above comparison further highlights the polarization of social interactions on the social media platform.

However, during the Italian elections considered in this work more than two political parties were present, so that the notion of polarization has to be extended to account for the presence of several opposing groups. For complex networks, a widely adopted approach is to use modularity maximization for group identification [25, 38, 39, 40]. Identified communities of users are characterized by intra-group connectivity denser than inter-group one.

In the case of the social bulk we can interpret modularity as an estimation of system's *social fragmentation* into more than two opposing groups. Here, we use the Louvain multilevel approach, known to be very efficient on large-scale networks [41]. In the whole social bulk we measure a fragmentation $\Phi_L = 0.812$, indicating the presence of several factions in the bulk network that are in a stronger opposition when compared to the null model ($\langle \Phi_L^{rand} \rangle = 0.692$, p-value $< 10^{-5}$). The fact that $\Phi_L > \Phi_F$ indicates that a more accurate description of the meso-scale organization of the social bulk is given when more than two groups are considered, in agreement with our hypothesis that results should reflect the real world socio-political scenario. To understand if this finding is robust or just an artefact due to how the social bulk is built, we have measured the social fragmentation of the original system during all phases of voting (see Fig 4). Once again, we observe that social fragmentation is stable across time and significantly larger than random expectation, confirming that results obtained from the social bulk are consistent.

However, neither polarization nor fragmentation can be used to quantify to which extent the system consists of isolated groups—the ones with no interactions with the rest of the system—which are effectively segregated in the network. Note that we are not referring to users in the periphery [42, 43] of the system, where information can slowly but flawlessly flow among all nodes in the network. Instead, we refer to groups unable to exchange information with the core of the system, i.e., to nodes belonging to disconnected components. We quantify *social segregation* Σ by considering the average number of connected components weighted by the number of their links (see Methods). If a social network consists of isolated nodes only, then $\Sigma = 1$, whereas $\Sigma = 0$ for systems with a single connected component. For a network consisting of M connected components of same size and density of interactions $\Sigma = 1 - 1/M$: The

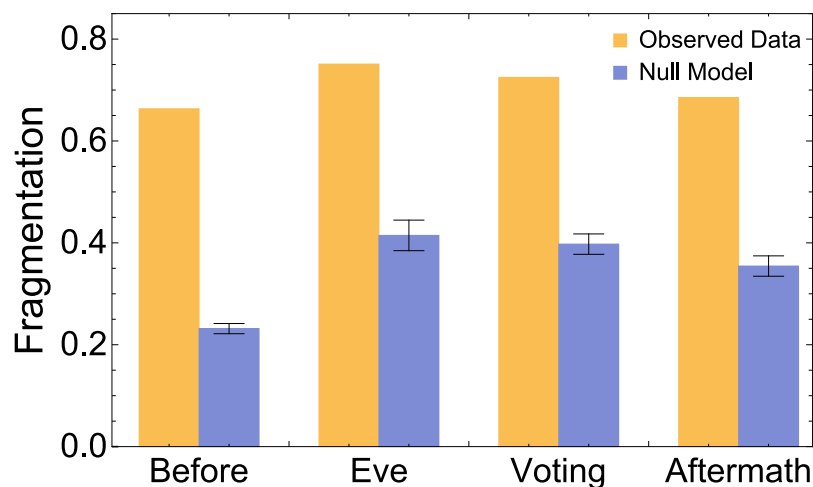


Fig 4. The online system is characterized by social fragmentation. **Top:** Fragmentation encodes the tendency of online users to organize in multiple opposing groups (see the text for further detail). During the four considered periods, the online social network is fragmented much more than random expectation. Small changes in fragmentation of the observed system across time are reflected in the null model, indicating that they can be explained by small changes in the heterogeneity of the underlying connectivity. Error bars indicate standard deviations.

<https://doi.org/10.1371/journal.pone.0214210.g004>

larger the number of components, the larger the social segregation. The segregation of the social bulk is $\Sigma = 0.476$, significantly stronger than random expectation ($\Sigma^{rand} = 0.172$ (p-value $< 10^{-5}$) based on a configuration model preserving the connectivity distribution. This indicates that strong interactions lead to more segregated components, with fewer bridges among connected components than expected from the heterogeneity of interactions only. Hence, the observed segregation represents additional evidence for the presence of antagonism in the considered social ecosystem.

Polarization, fragmentation and segregation analyses all constitute evidence that the social bulk displays densely connected groups in opposition with each other.

Groups in the social bulk highlight digital augmentation

Through the multilevel approach, we identified 8 main opposing communities (i.e. having more than 2% of the total nodes in the network), as reported in Table 2. The analysis of hubs in each group of the social bulk indicates that i) one group corresponds to a single augmented human and his/her bots; ii) five groups directly map the ecosystems of the main Italian political parties; iii) two groups encode news media universe, either traditional or online news organisations.

In this context, we provide an operative definition of augmented humans as human users having at least 50% of bot neighbours in the social bulk. Users with less than 3 bulk interactions are discarded. We systematically identified 1,010 user accounts (12.7% of humans in the social bulk) corresponding to augmented humans. The most central augmented human in terms of number of social interactions is User01 which interacts with 2,700 bots and 55 humans in the social bulk. We have anonymized the username for privacy purposes.

It is natural to wonder about how bots, humans and augmented humans are organized into communities within the social bulk. In fact, given the relevance of the voting event in the real world, our hypothesis is that communities should reflect real political movements and groups, to some extent.

Table 2. Largest online social groups. Most populated communities (with more than 250 users) in the social bulk, with top influencers listed per group. Top influencers are identified as hubs in the bulk network. As evident from the similarities among top influencers, groups reflect specific ecosystems of the Italian voting event: “Movimento 5 Stelle” (M5S), traditional media (Media), media with massive online presence (Web Media), “Partito Democratico” (PD), “Liberi e Uguali” (LEU), “Forza Italia” (FI), “Lega” and “Fratelli d’Italia” (Lega and FdI), and then the augmented human with all his/her interacting bots (Augmented Human). Bot (augmented) infiltration indicates the percentage of bot users (augmented humans) in each group. Excluding the community corresponding to the augmented human (made for 97.9% of bots), the mean bot infiltration in the bulk network is 29.2% while the mean augmented infiltration is 15.7%. The media groups are richer in bots as expected, since they include news media and online accounts of news papers. Note that users not corresponding to public groups, public entities or individuals with a public political profile (e.g., elected for a specific political party) have been anonymized. Interestingly, the account of User09 has been later suspended by Twitter for violating its policies.

Bulk Group	Top Influencers	Social Ecosystem	Bot infiltration	Augmented infiltration
C1	fattoquotidiano, User02, Mov5Stelle, User10, M5S_Europa, Puglia_M5S, User04	M5S	13.4%	13.3%
C2	repubblica, Agenzia_Ansa, sole24ore, TgLa7, RaiNews, Viminale, RTL1025, Corriere, MediasetTgcom24	Media	24.5%	12.1%
C3	SkyTG24, you_trend, nonleggerlo, ilpost, User05, User03, agorarai, espressoline	Web Media	26.6%	15%
C4	matteoreni, pdnetwork, CarloCalenda, User08, User09, PaoloGentiloni, PdMilano	PD	15.5%	15.4%
C5	liberi_uguali, civati, PietroGrasso, lauraboldrini, LiberiEugualiIT, SI_sinistra	LEU	16.3%	19.1%
C6	berlusconi, forza_italia, Elezioni2018_FI, renatobrunetta, User07, GruppoFICamera	FI	28.2%	22.1%
C7	User06, mattosalvinimi, borghi_claudio, AlbertoBagnai, GiorgiaMeloni, FratellidItalia, LegaSalvini	Lega and FdI	9.5%	12.9%
C8	User01	Augmented Human	97.9%	$3 \cdot 10^{-4}$ %

<https://doi.org/10.1371/journal.pone.0214210.t002>

First, we focused our attention on the augmented human's group, consisting of more than 2,500 automated agents artificially interacting with the augmented human user. This peculiar activity leads to a star-like structure for the corresponding community, as shown in Fig 3c, network C8. This finding has triggered our attention, driving our efforts towards quantifying the infiltration \mathcal{I}_s of a specific class s of accounts in each group, by considering the corresponding fraction of users in a given group (see Methods). Table 2 reports the infiltration of augmented humans in the groups of the social bulk. Unsurprisingly, infiltration of bots is higher in the group representing the augmented human and his/her automatic entourage of interacting social bots. Furthermore, we find that groups relative to news media are richer in bots compared to groups representing political parties, which is compatible with our previous finding of bots being preferentially news media broadcaster in the observed data.

The infiltration of augmented humans is approximately uniformly distributed across all identified groups, with the remarkable exception of C8, the augmented human's community. One would expect for the groups richer in bots to have also more augmented humans. Instead, bot and augmented human infiltration do not correlate with each other (Kendall Tau 0.07, p -value 0.8), indicating that augmented humans tend to interact selectively with the bots available in their groups rather than creating more. This trend is not valid for the group C8, where one human (User01) interacts almost exclusively with bots.

Testing the role of news media

In the analysis of the social bulk we identified two communities corresponding to news media accounts. In order to test for the influence of these information hubs on human-bot interactions, we performed a test in which we checked the robustness of our results when all users in the above two communities identifying news media accounts were not considered. The removal of news media accounts led to negligible fluctuations (around 0.02) in the fractions of human-bot interactions (cfr. Fig 1(a)) and in the total volume of tweets produced by bots (around 0.4%). These results indicate that a prominent amount of human-bot interactions does not involve news media accounts and it is not influenced by the presence of information hubs.

Augmented humans are hidden influencers

All the augmented humans identified in this study have, on average, less than 9,000 followers and 1,500 friends, indicating that a considerable amount of social influence was obtained by users that preferentially interacted with bots during the considered event. The analysis of information cascades revealed that almost 2 out of 3 augmented humans played an important role in the flow of online content: 67% of this class of users were either influentials or hidden influentials or broadcasters. Hidden influentials, known to be efficient spreaders in viral phenomena [44], are mostly humans but augmented humans also falls in this category (e.g. User01).

Groups in the social bulk reflect electoral outcomes

In order to investigate the representativeness of online groups in terms of real-world events beyond the hub analysis, we focused on the structural features of groups, namely the interaction volume of a group (i.e. the number of strong social interactions among users in the group) and the group size (i.e. the number of users in a given group). In Table 3, we show that the outcome of Italian elections (i.e. the fraction of votes received by each political group) strongly correlates with the group volume (Spearman rank correlation coefficient $\rho = 0.9$, p -value = 0.039). This correlation is statistically significant within a 5% significance level and

Table 3. Network analysis of groups in the social bulk reflect election outcomes. The five political ecosystems from the bulk network are ranked against their topological features: i) interaction volume, i.e. the number of social actions within the group; ii) size, i.e. the number of individuals in the group. The rank based on online interactions strongly mirrors the election outcome (Spearman $\rho = 0.9$, p-value = 0.039), supporting the hypothesis that online social interactions are tightly entwined to outcomes and events in the real-world.

Group	Election Outcome (% votes)	Rank	Interaction Volume	Rank	Group Size	Rank
M5S	32.68%	1	3,857 (34.82%)	1	1,162	2
Lega and FdI	20.76%	2	3,205 (28.93%)	2	1,491	1
PD	18.72%	3	2,504 (22.60%)	3	1,133	3
FI	14.01%	4	675 (6.09%)	5	330	5
LEU	3.39%	5	837 (7.55%)	4	434	4

<https://doi.org/10.1371/journal.pone.0214210.t003>

direct sampling of rankings was used in order to compute the p-value without relying on any assumption about the large-scale statistical properties of the data. The strong correlation found indicates that the volume of online interactions closely mirrors the election outcome for this case study, although further research is needed to generalize this result and confirm its predictive power.

Discussion

Online social systems and the information they continuously generate provide an invaluable resource for computational social scientists and their large-scale analysis of human behavior [45, 36, 46] and the emergence of collective attention [2, 47]. The analysis of information and behavioral spreading on social media [17] revealed that an individual is much more likely to adopt a content when his/her neighbors in the social network tend to reinforce it [48]. On the one hand, this allows online media to facilitate, for instance, the dissemination of emergency information and help coordinate relief efforts [49]. On the other hand, the same social networks can be misused to spread fake content farther, faster and deeper [13].

In this work we have identified and quantified a new phenomenon, i.e. digital augmentation, to characterize individuals that coordinate from hundreds to thousands of social bots for achieving a social influence comparable to the one of political parties and news media organizations, with serious repercussions in the real-world.

Our results strongly support the idea that via augmentation even common users can become social influencers without having a large social neighborhood but rather by recurring to the aid of either armies of bots or the selection of a few key helping bots. This digital augmentation represents an interesting behavioral response aimed at overcoming the well documented pressure for achieving influence and recognition in online ecosystems [16, 4, 6, 31] and during voting events [19]. While in real life such augmentation comes mainly from smart devices, our work presents compelling evidence that in online social platforms the augmentation for achieving social influence is represented by an exploitation of social bots by human accounts.

Furthermore, the strong correlation between the volume of online interactions in the social bulks and the electoral outcomes highlights the role potentially played by online social systems during the voting process. This finding is in full agreement with previous works showing how online ecosystems acted upon society by altering the emotions [19] and beliefs [3, 4] of large populations of individuals. It is worth underlining that the observed groups are relative to the network structure of social endorsements: Considering the layout of online endorsement can provide information beneficial for more accurate predictions of electoral outcomes. Further investigation of online social systems under the perspective of predicting electoral outcomes would provide interesting challenges for future work.

Our work provides a first step towards a more systematic quantification of the impact of digital augmentation in opinion formation and the manipulation of online attention by means of human-bot interactions.

Supporting information

S1 Text. Supplementary text and figures.

(PDF)

Acknowledgments

We acknowledge Pierluigi Sacco for insightful discussion.

Author Contributions

Conceptualization: Manlio De Domenico.

Data curation: Manlio De Domenico.

Formal analysis: Massimo Stella.

Methodology: Massimo Stella, Marco Cristoforetti, Manlio De Domenico.

Software: Marco Cristoforetti.

Supervision: Manlio De Domenico.

Validation: Massimo Stella.

Visualization: Manlio De Domenico.

Writing – original draft: Massimo Stella, Marco Cristoforetti, Manlio De Domenico.

Writing – review & editing: Massimo Stella, Marco Cristoforetti, Manlio De Domenico.

References

1. De Domenico M, Lima A, Mougél P, Musolesi M. The anatomy of a scientific rumor. *Scientific reports*. 2013; 3:2980. <https://doi.org/10.1038/srep02980> PMID: 24135961
2. Borge-Holthoefer J, Perra N, Gonçalves B, González-Bailón S, Arenas A, Moreno Y, et al. The dynamics of information-driven coordination phenomena: A transfer entropy analysis. *Science advances*. 2016; 2(4):e1501158. <https://doi.org/10.1126/sciadv.1501158> PMID: 27051875
3. Muchnik L, Aral S, Taylor SJ. Social influence bias: A randomized experiment. *Science*. 2013; 341(6146):647–651. <https://doi.org/10.1126/science.1240466> PMID: 23929980
4. Bond RM, Fariss CJ, Jones JJ, Kramer AD, Marlow C, Settle JE, et al. A 61-million-person experiment in social influence and political mobilization. *Nature*. 2012; 489(7415):295–298. <https://doi.org/10.1038/nature11421> PMID: 22972300
5. Wagner C, Mitter S, Körner C, Strohmaier M. When social bots attack: Modeling susceptibility of users in online social networks. *Making Sense of Microposts (# MSM2012)*. 2012; 2(4):1951–1959.
6. Cresci S, Di Pietro R, Petrocchi M, Spognardi A, Tesconi M. Fame for sale: efficient detection of fake Twitter followers. *Decision Support Systems*. 2015; 80:56–71. <https://doi.org/10.1016/j.dss.2015.09.003>
7. Cresci S, Di Pietro R, Petrocchi M, Spognardi A, Tesconi M. The paradigm-shift of social spambots: Evidence, theories, and tools for the arms race. In: *Proceedings of the 26th International Conference on World Wide Web Companion*. International World Wide Web Conferences Steering Committee; 2017. p. 963–972.
8. Gilani Z, Kochmar E, Crowcroft J. Classification of Twitter Accounts into Automated Agents and Human Users. In: *Proceedings of the 2017 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2017*. ACM; 2017. p. 489–496.

9. Varol O, Ferrara E, Davis CA, Menczer F, Flammini A. Online human-bot interactions: Detection, estimation, and characterization. arXiv preprint arXiv:170303107. 2017;.
10. Ferrara E, Varol O, Davis C, Menczer F, Flammini A. The rise of social bots. *Comm of the ACM*. 2016; 59(7):96–104. <https://doi.org/10.1145/2818717>
11. Ferrara E. Disinformation and social bot operations in the run up to the 2017 French presidential election. *First Monday*. 2017; 22(8). <https://doi.org/10.5210/fm.v22i8.8005>
12. Stella M, Ferrara E, Domenico MD. Bots increase exposure to negative and inflammatory content in online social systems. *Proceedings of the National Academy of Sciences*. 2018; p. 201803470.
13. Vosoughi S, Roy D, Aral S. The spread of true and false news online. *Science*. 2018; 359(6380):1146–1151. <https://doi.org/10.1126/science.aap9559> PMID: 29590045
14. González-Bailón S, Borge-Holthoefer J, Rivero A, Moreno Y. The dynamics of protest recruitment through an online network. *Scientific reports*. 2011; 1:197. <https://doi.org/10.1038/srep00197> PMID: 22355712
15. González-Bailón S, Borge-Holthoefer J, Moreno Y. Broadcasters and hidden influentials in online protest diffusion. *American Behavioral Scientist*. 2013; 57(7):943–965. <https://doi.org/10.1177/0002764213479371>
16. Aral S, Muchnik L, Sundararajan A. Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks. *Proceedings of the National Academy of Sciences*. 2009; 106(51):21544–21549. <https://doi.org/10.1073/pnas.0908800106>
17. Aral S, Walker D. Identifying influential and susceptible members of social networks. *Science*. 2012; 337(6092):337–341. <https://doi.org/10.1126/science.1215842> PMID: 22722253
18. Ashworth S, Bueno de Mesquita E, Friedenberg A. Learning about voter rationality. *American Journal of Political Science*. 2018; 62(1):37–54. <https://doi.org/10.1111/ajps.12334>
19. Bruter M, Harrison S. Understanding the emotional act of voting. *Nature Human Behaviour*. 2017; 1(0024):1–3.
20. He W, Zha S, Li L. Social media competitive analysis and text mining: A case study in the pizza industry. *International Journal of Information Management*. 2013; 33(3):464–472. <https://doi.org/10.1016/j.ijinfomgt.2013.01.001>
21. Emmert-Streib F, Yli-Harja OP, Dehmer M. Data analytics applications for streaming data from social media: What to predict? *Frontiers in Big Data*. 2018; 1:2. <https://doi.org/10.3389/fdata.2018.00002>
22. Chu Z, Gianvecchio S, Wang H, Jajodia S. Who is tweeting on Twitter: human, bot, or cyborg? In: *Proceedings of the 26th annual computer security applications conference*. ACM; 2010. p. 21–30.
23. Savulescu J, Bostrom N. *Human enhancement*. Oxford University Press on Demand; 2009.
24. Ding CH, He X, Zha H, Gu M, Simon HD. A min-max cut algorithm for graph partitioning and data clustering. In: *Data Mining, 2001. ICDM 2001, Proceedings IEEE International Conference on*. IEEE; 2001. p. 107–114.
25. Newman ME. Modularity and community structure in networks. *Proceedings of the national academy of sciences*. 2006; 103(23):8577–8582. <https://doi.org/10.1073/pnas.0601602103>
26. Metaxas PT, Mustafaraj E, Wong K, Zeng L, O’Keefe M, Finn S. What Do Retweets Indicate? Results from User Survey and Meta-Review of Research. In: *ICWSM*; 2015. p. 658–661.
27. Watts DJ. A simple model of global cascades on random networks. *Proceedings of the National Academy of Sciences*. 2002; 99(9):5766–5771. <https://doi.org/10.1073/pnas.082090499>
28. Gleeson JP, Durrett R. Temporal profiles of avalanches on networks. *Nature Communications*. 2017; 8(1):1227. <https://doi.org/10.1038/s41467-017-01212-0> PMID: 29089481
29. Goel S, Watts DJ, Goldstein DG. The structure of online diffusion networks. In: *Proceedings of the 13th ACM conference on electronic commerce*. ACM; 2012. p. 623–638.
30. Martin T, Hofman JM, Sharma A, Anderson A, Watts DJ. Exploring limits to prediction in complex social systems. In: *Proceedings of the 25th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee; 2016. p. 683–694.
31. Gleeson JP, O’Sullivan KP, Baños RA, Moreno Y. Effects of network structure, competition and memory time on social spreading phenomena. *Physical Review X*. 2016; 6(2):021019. <https://doi.org/10.1103/PhysRevX.6.021019>
32. Dorogovtsev SN, Goltsev AV, Mendes JFF. Critical phenomena in complex networks. *Rev Mod Phys*. 2008; 80:1275–1335. <https://doi.org/10.1103/RevModPhys.80.1275>
33. Ferrara E. Contagion dynamics of extremist propaganda in social networks. *Information Sciences*. 2017; 418:1–12. <https://doi.org/10.1016/j.ins.2017.07.030>
34. Bastos MT, Mercea D. Serial activists: Political Twitter beyond influentials and the twittertariat. *New Media & Society*. 2016; 18(10):2359–2378. <https://doi.org/10.1177/1461444815584764>

35. Bakshy E, Hofman JM, Mason WA, Watts DJ. Everyone's an influencer: quantifying influence on twitter. In: Proceedings of the fourth ACM international conference on Web search and data mining. ACM; 2011. p. 65–74.
36. Conover M, Ratkiewicz J, Francisco MR, Gonçalves B, Menczer F, Flammini A. Political polarization on twitter. ICWSM. 2011; 133:89–96.
37. Newman M. Networks: an introduction. Oxford university press; 2010.
38. Fortunato S. Community detection in graphs. Physics reports. 2010; 486(3-5):75–174. <https://doi.org/10.1016/j.physrep.2009.11.002>
39. Fortunato S, Hric D. Community detection in networks: A user guide. Physics Reports. 2016; 659:1–44. <https://doi.org/10.1016/j.physrep.2016.09.002>
40. Newman ME. Communities, modules and large-scale structure in networks. Nature physics. 2012; 8(1):25. <https://doi.org/10.1038/nphys2162>
41. Blondel VD, Guillaume JL, Lambiotte R, Lefebvre E. Fast unfolding of communities in large networks. Journal of statistical mechanics. 2008; 2008(10):P10008. <https://doi.org/10.1088/1742-5468/2008/10/P10008>
42. Borgatti SP, Everett MG. Models of core/periphery structures. Social networks. 2000; 21(4):375–395. [https://doi.org/10.1016/S0378-8733\(99\)00019-2](https://doi.org/10.1016/S0378-8733(99)00019-2)
43. Rombach P, Porter MA, Fowler JH, Mucha PJ. Core-periphery structure in networks (revisited). SIAM Review. 2017; 59(3):619–646. <https://doi.org/10.1137/17M1130046>
44. Baños RA, Borge-Holthoefer J, Moreno Y. The role of hidden influentials in the diffusion of online information cascades. EPJ Data Science. 2013; 2(1):6. <https://doi.org/10.1140/epjds18>
45. Lazer D, Pentland AS, Adamic L, Aral S, Barabasi AL, Brewer D, et al. Life in the network: the coming age of computational social science. Science (New York, NY). 2009; 323(5915):721. <https://doi.org/10.1126/science.1167742>
46. Ruths D, Pfeffer J. Social media for large studies of behavior. Science. 2014; 346(6213):1063–1064. <https://doi.org/10.1126/science.346.6213.1063> PMID: 25430759
47. Baronchelli A. The emergence of consensus: a primer. Royal Society open science. 2018; 5(2):172189. <https://doi.org/10.1098/rsos.172189> PMID: 29515905
48. Centola D. The spread of behavior in an online social network experiment. Science. 2010; 329(5996):1194–1197. <https://doi.org/10.1126/science.1185231> PMID: 20813952
49. Kryvasheyev Y, Chen H, Obradovich N, Moro E, Van Hentenryck P, Fowler J, et al. Rapid assessment of disaster damage using social media activity. Science Advances. 2016; 2(3):e1500779. <https://doi.org/10.1126/sciadv.1500779> PMID: 27034978