

RESEARCH

Open Access



Bot stamina: examining the influence and staying power of bots in online social networks

Ross Schuchard^{1*} , Andrew T. Crooks^{1,2}, Anthony Stefanidis^{2,3} and Arie Croitoru³

* Correspondence: rschucha@gmu.edu

¹Computational Social Science Program, Department of Computational and Data Sciences, George Mason University, Fairfax, Virginia 22030, USA
Full list of author information is available at the end of the article

Abstract

This study presents a novel approach to expand the emergent area of social bot research. We employ a methodological framework that aggregates and fuses data from multiple global Twitter conversations with an available bot detection platform and ultimately classifies the relative importance and persistence of social bots in online social networks (OSNs). In testing this methodology across three major global event OSN conversations in 2016, we confirmed the hyper-social nature of bots: suspected social bot accounts make far more attempts on average than social media accounts attributed to human users to initiate contact with other accounts via retweets. Social network analysis centrality measurements discover that social bots, while comprising less than 0.3% of the total corpus user population, display a disproportionately high level of structural network influence by ranking particularly high among the top users across multiple centrality measures within the OSN conversations of interest. Further, we show that social bots exhibit temporal persistence in centrality ranking density when examining these same OSN conversations over time.

Keywords: Social bot analysis, Computational social science, Social network analysis, Online social networks

Introduction

As online social network (OSN) platforms (e.g. Twitter, Instagram, Sina Weibo) continue to attract dramatic global participation in terms of active user rates, they are becoming indispensable components of the online ecosystem (Blackwell et al. 2017). In the same sense that Fuchs (2005) describes the Internet as a socio-technological system, user devotion to OSNs has led usage patterns that transcend simple messaging activities among networks of friends. In the United States (U.S.), OSN platforms recently surpassed print newspapers as a primary source for news, and they continue to gain traction in relation to other traditional news sources such as television and radio (Mitchell 2018). While the convenience of receiving ‘news’ within a multipurpose communication system is understandable, the sharing of real-world news in a social interaction environment may lead to unintended consequences. Sunstein (2018) suggests that the homophily-driven nature of OSNs results in the formation of echo

chambers, which serve as fertile ground for the amplification of perpetuated false information, or fake news, among their members.

Recent studies have pointed to evidence of fake news within OSN conversations (e.g. Lazer et al. 2018; Grinberg et al. 2019). Furthermore, while examining news stories within Twitter from 2006 to 2017, Vosoughi et al. (2018) discovered that false stories spread more rapidly and to a greater audience than true stories. In addition to struggling to decipher the veracity of news, OSNs also have trouble accounting for the veracity of user accounts. This is largely due to the proliferation of accounts belonging to social bots, which are computer algorithms designed to mimic human behavior and interact with humans in an automated fashion. While automated in nature, social bots are not universally designed for intentional malice, as many bots serve in benign or even helpful roles (e.g. news aggregator) (Ferrara et al. 2016). The increasing sophistication of bots has made it difficult for human users to discern fellow human users from social bots in OSNs (Ruths and Pfeffer 2014; Ferrara et al. 2016). While Vosoughi et al. (2018) argued that social bots were responsible for spreading both false and true news at the same rates as humans, Shao et al. (2018) discovered that social bots amplified news stories from low-credible sources in a disproportionate fashion. Although such studies have demonstrated the strong presence of social bots in OSNs, the full extent to which these bots introduce, spread or amplify information remains elusive. For this reason, it is essential to gain greater understanding of the implications associated with human and machine dialogue, either intentional or not.

Initial social bot research continues to build upon a foundation of the classification and detection of social bots in OSNs (Chu et al. 2012; Davis et al. 2016; Chavoshi et al. 2016). The increasing sophistication of bots and the ability of some bots to mimic human behavior are proving to be too complex for current passive detection methods (Cresci et al. 2017). Even some simple rules-based social bots continue to gain an influential role in networks and go undetected for extended periods of time (Abokhodair et al. 2015). Recent promising advances in active bot detection algorithm development follow an adversarial learning approach by employing genetic algorithms to detect evolving bots (Cresci et al. 2018b, 2019b, 2019c). While bot detection methodologies are improving with respect to keeping pace with evolving bot sophistication, there exists ample opportunities to develop and test necessary social bot analysis techniques to better characterize currently detectable social bots. Recent initial social bot analysis studies, which rely upon an array of multidisciplinary approaches, have provided positive insights into social bot influence within OSN conversations involving healthcare issues (Broniatowski et al. 2018), elections (Howard et al. 2018; Stella et al. 2018), financial trading markets (Cresci et al. 2018a, 2019a) and protests (Suárez-Serrato et al. 2016). Given that social bots aim to mimic and replicate human behavior, some researchers suggest that a computational social science (CSS) paradigm could provide a compelling framework for characterizing the influence that bots may have on OSN conversations (Ciampaglia 2018; Strohmaier and Wagner 2014).

It is from a CSS perspective that we present a unique methodology and analysis framework to observe human and social bot behavior and interactions within OSN conversations. Specifically, we acquire Twitter data associated with three major global events in 2016: the 2016 U.S. presidential election primary races, the ongoing Ukrainian conflict involving Russia and Ukraine, and the Turkish

government's implementation of censorship practices against its own citizens. We then enrich the Twitter data by classifying the bot status of all user accounts within the corpus. This enables a multi-faceted data analysis approach that includes comparative descriptive statistical analysis methods and social network analysis techniques to determine the relative importance and persistence of social bots within each global conversation. Overall, we construct a corpus consisting of over 28.6 million tweets produced by approximately 5 million distinct users, of which, we label 14,386 of those users as likely social bots producing more than 1.9 million tweets. This reproducible framework, which can be extended to other OSN conversations and additional bot detection algorithms, creates an opportunity to better describe currently detected bots, while also providing essential feedback loops to bot detection research.

The results of this study show that suspected social bot users, on average, attempt to initiate contact with other users via retweets at a rate far higher than human users. Through the application of social network analysis centrality measurements, we discover that social bots, while comprising less than 0.3% of the total user population, display a profound level of structural network influence by ranking particularly high among the top eigenvector centrality users within the U.S. Election and the Ukraine Conflict OSN conversations. Further, in observing the temporal persistence of suspected social bots, we find that bot users maintain their density of top centrality rankings over the cumulative OSN conversations of interest. Finally, the most relatively influential social bots from our Twitter corpus display a distinct ability to attract higher in-degree edge connections from human users that retweet their original bot messages. These results are quite promising given this study relied upon one open-source bot detection platform which provides limited total conversational coverage, but precise positive bot classification.

In an earlier paper (Schuchard et al. 2019) we presented a study that focused on developing an initial framework to characterize the pervasiveness and relative importance of social bots in OSNs. This current paper extends this earlier work by providing a more robust contribution along three lines of effort. First, we expand our analysis to include additional centrality measures that are specific to complex communicative networks. Second, we develop temporal centrality rank persistence results for each online conversation to determine the relative staying power of certain social bots over time. Finally, we examine the evolution of ego networks for the most structurally relevant bots over time in an effort to better characterize the user types communicating with social bots.

The remainder of this paper is structured as follows. In the “[Background](#)” section, a brief synopsis introduces current social bot detection methods and social bot analysis efforts. “[Enabling a social bot analysis framework](#)” provides a detailed overview of this study's processes, which acquire and fuse the data sources to enable the subsequent analysis section. “[Analysis results and discussion](#)” focuses on the results of the comparative descriptive statistical analysis methods and social network analysis techniques from this study and discusses their implications across the global event use-cases of interest. Finally, we conclude in “[Conclusion](#)” section and highlight potential future research opportunities.

Background

In the past 15 years, the digital data exhaust created by increasing OSN usage rates and the relative ease at which researchers can gain access to such data has led to the rapid emergence of social media research. As social media research norms continue to develop, studies have produced insights from OSN-extracted data on topics including disaster response (Crooks et al. 2013; Sakaki et al. 2013; Avvenuti et al. 2016a, 2016b), mental illness forecasting (Reece et al. 2017) and political polarization (Conover et al. 2011). The limitations and risks associated with using OSN data for research are well documented (Tufekci 2014; Ruths and Pfeffer 2014), but the adaptive nature of social bots participating in OSNs amplify these concerns and may lead to many additional research implications (Morstatter et al. 2016).

The evidentiary rise of social bots in OSNs has led to a corresponding increase in research dedicated to bot detection (Murthy et al. 2016). The motivation and design methods associated with bots can vary dramatically, so a myriad of detection methods is necessary to account for the potential characteristics or activities attributable to certain social bots. In the following, we focus on two bot detection platforms, Botometer (Davis et al. 2016) and DeBot (Chavoshi et al. 2016), which exhibit very dissimilar design criteria but are both widely used in research due to the fact that they provide open access through web applications and application programming interfaces (APIs).

The Botometer (formerly named BotOrNot) bot detection platform employs a supervised ensemble Random Forest classification technique, which classifies potential Twitter accounts as bots according to six different classifiers based on more than 1,000 extracted features from an associated Twitter account (Davis et al. 2016). Botometer assigns a probabilistic [0,1] score representing the likelihood that a Twitter account is a bot, with simple and sophisticated bots falling within score ranges of 0.8–1.0 and 0.5–0.7, respectively (Varol et al. 2017). The DeBot bot detection platform, on the other hand, relies upon an unsupervised warped correlation method to find correlated Twitter accounts that have more than 40 synchronous events within a given time window (Chavoshi et al. 2016). DeBot provides a binary positive or negative bot classification for a Twitter account at incredibly high levels of precision (Chavoshi et al. 2017), but at a cost of recall due to evaluating smaller populations of Twitter accounts (Morstatter et al. 2016). In contrast to Botometer, DeBot archives its detection results, which allows researchers to ascertain potential bot status for previously detected accounts from a historical perspective (Chavoshi et al. 2017). Botometer provides a bot evaluation score based on the time of a given query and does not provide a retrospective analysis capability. As Cresci et al. (2017) aptly asserts, individual bot detection methodologies are not designed to detect the wide range of operational social bot types, and they require continual refinement to keep pace with evolving bot sophistication.

Social bot analysis is gaining traction as a means to better understand the impact of social bots and potentially provide essential feedback to bot detection research efforts. While social bot analysis currently lacks a formal definition, we submit an informal definition to be a multidisciplinary research effort employing quantitative and/or qualitative methods with a stated purpose of better understanding detectable social bots and their behaviors in OSNs. Recent initial social bot analysis contributions examine the presence of detected social bots in Twitter conversations involving the 2016 U.S. presidential election (Bessi and Ferrara 2016; Howard et al. 2018), the United Kingdom

Brexit referendum (Howard and Kollanyi 2016; Duh et al. 2018), the ongoing Ukrainian-Russian conflict (Hegelich and Janetzko 2016), financial trading markets (Cresci et al. 2018a, 2019a) and the debates on vaccination (Broniatowski et al. 2018). These works have built the initial corpus of social bot analysis research, but much work is left to be done to introduce more advanced evaluation methods across greater use-cases of interest. One path to advancing these evaluation methods are social network analysis (SNA) techniques.

Observable human and bot interactions in OSN platforms such as Twitter provide a prime opportunity to employ SNA techniques to evaluate the relative importance of detected bots in comparison to human users. A key finding in Boshmaf et al. (2013), Aiello et al. (2014) and Mønsted et al. (2017) is that social bot infiltration and subsequent interactions with human users in OSNs occur at surprisingly high rates. Learning from Cha et al. (2010) that relative influence in Twitter by users is not necessarily gained through popularity (i.e. associated follower volume), we can look to SNA techniques to derive influence in OSNs (Kwak et al. 2010; Weng et al. 2010; Bakshy et al. 2011; Riquelme and González-Cantergiani 2016).

Initial social bot research employing advanced social network analysis techniques to evaluate bot influence in OSNs is limited but growing. Aiello et al. (2014) applies the PageRank and Hypertext Induced Topic Search (HITS) link analysis algorithms to judge the relative importance of an experimental bot. In observing the Catalan referendum Twitter conversation, Stella et al. (2018) uses an average PageRank valuation to compare suspected bot and human accounts, while also showing that bot interactions targeting human accounts positively correlates with the in-degree of human-to-human interactions. Perna and Tagarelli (2018) present the most promising effort to quantify social bot relevance with their ensemble machine learning framework, Learning-To-Rank-Social-Bots (LTRSB). The LTRSB framework aims to provide a unifying method to rank bots based on the extracted features present in the available bot detection platforms (e.g. Botometer, DeBot, BotWalk) (Perna and Tagarelli 2018).

OSN research has turned into a burgeoning field in academia that has risen in stride with the overall rapid advancement in global social media usage. The increasing reliance upon OSN platforms as primary news sources by today's digitally-focused citizens, however, highlights the need to better identify and analyze the implications of social bot actors participating in online dialogue. The research presented in this paper adds to the field of social bot analysis by introducing a novel methodology to ascertain the relative importance and persistence of social bots across multiple OSN conversation use cases. This reproducible methodology seeks to enable a comparative framework to extend to other OSN conversation use cases, while also accounting for the incorporation of future bot detection algorithms.

Enabling a social bot analysis framework

This study develops and employs a social bot analysis framework that focuses on the aggregation of multiple harvested Twitter conversations and bot detection results to better characterize the relative influence and persistence of social bots in OSNs. It is in this section that we describe the processes to transform these data, which enable the ensuing ensemble application of comparative descriptive analysis methods and SNA

techniques in this study. Figure 1 summarizes the processes comprising our social bot analysis framework and detailed subsections constitute the remainder of this section. “Data acquisition and processing” presents the details behind each OSN conversation of interest and the associated keywords serving as the input parameters to harvest tweets. “Bot enrichment” describes the bot labeling process for each Twitter user in this study’s corpus. “Retweet network construction” explains the process to build network objects for each OSN conversation, while “Data analysis” concludes the section by introducing the analysis methods comprising the subsequent sections of the study.

Data acquisition and processing

Three major global events from 2016 serve as the OSN conversation use cases in this study. Focusing solely on Twitter, we examine harvested tweets from 4 weeks of different topical conversations, to include a political election (2016 U.S. Presidential Election), a war/conflict (2016 Ukraine Conflict) and censorship (2016 Turkish Censorship). By analyzing varied topics, we seek to determine social bot behavioral differences across a diverse set of conversations. We introduce and summarize briefly the three OSN conversations in what follows:

U.S. presidential election (February 1–28, 2016)

We observe 4 weeks of tweets in February 2016 based on keywords associated with the 2016 U.S. Presidential Election. During this period, the election’s primary races to determine the Republican and Democratic party candidates for the general election are well underway. The Republican primary race attracts considerable social media attention as then-candidate Donald Trump gains substantial momentum towards securing the Republican nomination over Texas Senator Ted Cruz. The Democratic race develops into a two-candidate race between former Secretary of State Hillary Clinton and Vermont Senator Bernie Sanders.

Ukraine conflict (August 1–28, 2016)

We observe 4 weeks of tweets in August 2016 based upon keywords associated with the ongoing conflict between Ukraine and Russia. At this point in time, it has been fewer than 3 years since the anti-Russian Euromaidan protests and the subsequent annexation of Crimea by Russia. Military bravado and political rhetoric between these

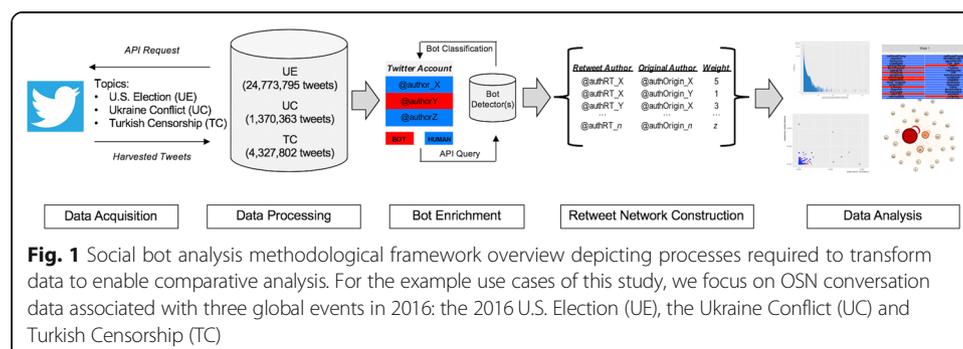


Fig. 1 Social bot analysis methodological framework overview depicting processes required to transform data to enable comparative analysis. For the example use cases of this study, we focus on OSN conversation data associated with three global events in 2016: the 2016 U.S. Election (UE), the Ukraine Conflict (UC) and Turkish Censorship (TC)

nations increases dramatically as the 25th anniversary of Ukrainian independence from Russia approaches (August 24, 1991).

Turkey censorship (December 1–28, 2016)

We observe 4 weeks of tweets based on keywords associated with Turkish government censorship of OSNs, specifically Twitter, in December 2016. Following a failed coup attempt against the sitting Turkish government in July 2016, government officials are keen to monitor and suppress messaging campaigns on OSNs. In December 2016, the Turkish government explicitly blocks Turkish citizens from using Twitter immediately following two events. The first block period takes place in the aftermath of the public assassination of Andrei Karlov, the Russian Ambassador to Turkey, on December 19, 2016. Turkey initiates a second block on December 23rd immediately following the release of a video that shows two Turkish soldiers being burned alive.

Based on the OSN conversation overviews as described above, we extract what we deem to be the representative keywords for each topic as shown in Table 1. We then submit these keywords to harvest associated tweets via the Twitter Standard Search API. Overall, our keyword harvest yields more than 28.6 million total tweets produced by approximately 5 million unique accounts with a breakdown for each OSN conversation as follows: U.S. Presidential Election ~ 23.3 million tweets (~ 3.3 million unique accounts), Ukraine Conflict ~ 1.3 million tweets (~ 0.4 million unique accounts) and Turkish Censorship ~ 4.0 million tweets (~ 1.3 million unique accounts). In order to account for the storage and computation demands for such a large data corpus, we conduct all processing of the data within an Amazon Web Services (AWS) EC2 t2.2xlarge instance consisting of 8 vCPUs and 32GiB of RAM. In doing so, we are able to rapidly create specified data objects for processing at the local level, while also maintaining a scalable compute/storage platform to account for future data expansion.

Bot enrichment

To individually label each unique Twitter account user in our tweet corpus as a human or a suspected bot, we employ the open-source DeBot bot detection platform (Chavoshi et al. 2016). DeBot was the logical bot detection platform to use as the detection service proof of concept for this study since, as the “Background” section details, the archival nature of DeBot allows us to classify our historical Twitter user accounts from 2016. Further, DeBot, via its unsupervised warped correlation method, detects bots at much higher precision rate than other bot detection platforms (Chavoshi et al. 2017). In addition to the historical limitation of using Botometer, there is an additional limitation restricting Botometer’s use for this study. Botometer underwent a major platform upgrade¹ in 2018, which resulted in a new scoring baseline that is not comparable to previous Botometer results. While such precision comes at the cost of lower recall and increases the risk of false-negative bots (i.e. automatically assessing non-assessed accounts as human accounts) as Morstatter et al. (2016) notes, we feel DeBot is the logical platform to initially test our social bot analysis framework given the historical

¹Further Botometer platform details, to include scoring algorithm changes, can be found at <https://botometer.iuni.iu.edu/#!/api-!/faq>.

Table 1 Keyword list of terms for submission to Twitter Standard Search API for each OSN conversation of interest in this study

U.S. Election ^a (Language)	Ukraine Conflict (Language)	Turkish Censorship (Language)
trump (English)	clinton (English)	turkey (English)
@realdonaldtrump (English)	hillary (English)	türkei (Turkish)
cruz (English)	sanders (English)	turkish (English)
@tedcruz (English)	bernie (English)	direntwitter (English)
makeamericagreatagain (English)	gop (English)	resisttwitter (English)
trump2016 (English)	gopdebate (English)	occupytwitter (English)
	ukraine (English)	erdoğanblockdwtwitter (English)
	ukrainian (English)	erdoğanblockstwitter (English)
	україна (Russian)	twitterisblockedinturkey (English)
	russia (English)	direntwitter (English)
	russian (English)	erdoğan (Turkish)
	kiev (English)	turkeycoup (English)
	киев (Russian)	
	київ (Ukrainian)	
	crimea (English)	
	spetsnaz (English)	
	specnaz (English)	
	putin (English)	

^a Keywords submitted for the 2016 U.S. Election skew heavily toward capturing tweets associated with then-candidate Donald Trump. We provide representative inclusion of other candidates but attempt to capitalize on the almost celebrity status of Mr. Trump by including additional Trump-related keywords

nature of our Twitter corpus. As we further stress in the “[Conclusion](#)” section of the paper, future improvements in social bot analysis research will rely upon the increased availability of additional bot detection algorithms to researchers which will allow for a more comprehensive coverage of all types of bots.

The bot enrichment process entails extracting all unique tweet author names from this study’s tweet corpus and passing them for classification via the DeBot API.² The returns simply classify the tweet author name as a suspected social bot or not (i.e. a human author). We then automatically label each user account through parsing scripts and merge the bot classification results with the tweet corpus. This process is easily extendible to account for other bot detection platform results. While beyond the scope of our study due to the historical nature of our tweet corpus, future work should also consider tracking the suspension/deletion statuses of accounts as the typical activities of social bot accounts make them primary targets of such actions by Twitter (Ferrara 2017).

We ultimately label 14,386 Twitter user accounts as likely social bots based on the DeBot classification results. This includes restricting positive bot labels to accounts only evaluated by DeBot prior through the dates of our Twitter corpus. While this population represents just 0.29% of the total unique user accounts in the corpus, social bots are very active, and account for an over twentyfold share of the corpus of published tweet (1,966,623 tweets, or 6.80% of the total) and over thirtyfold share of the corpus of published retweets (1,495,388, or 8.84% of the total). Table 2 below provides weekly and cumulative corpus metrics for each of the OSN conversations of interest. At the specific OSN conversation level, the U.S. Election corpus shows much greater weekly and cumulative tweet and retweet percentage contributions from social bot user accounts in comparison to the Ukraine Conflict and Turkey Censorship corpuses, even though the relative percentage of total bot accounts is much smaller in the election corpus. Further, social bots account for a higher percentage of total retweets in comparison to tweets across all conversations.

Retweet network construction

The practice of retweeting can produce a diverse range of conversational implications, but Twitter users that deliberately retweet are more likely trying to engage in conversation or directly share information (Boyd et al. 2010). In our study, retweets account for 58.4% (~ 16.7 million) of the total tweet corpus, with the specific conversation retweet densities of 57.6%, 50.3% and 65.9% for the U.S. Election, the Ukraine Conflict and the Turkish Censorship conversations, respectively. The act of a retweet between two Twitter users (i.e. nodes) results in an observable directed network connection (i.e. an edge). We assign a directed edge weight value of ‘1’ for each initial directed retweet connection between two users and increase the edge weight for each additional number of retweets between the appropriate directional pair of users.

A retweet serves as the primary artifact from which we can extract a ‘node-edge’ network construct from a Twitter conversation and ultimately enables the application of the SNA methods we introduce in the subsequent “[Data analysis](#)” section. In total, each

²The DeBot API is accessible at <https://www.cs.unm.edu/~chavoshi/debot/api.html>.

Table 2 Twitter corpus overview at the weekly and cumulative perspective for each OSN conversation in this study

Corpus	Week(s)	Tweets	Retweets	Users
United States Election Bot Source (% of total)	Week 1	4,054,560	2,280,176	1,029,090
	(Feb. 1–7, 2016)	315,540 (7.78%)	235,815 (10.34%)	4,229 (0.41%)
	Week 2	4,991,968	2,802,381	997,107
	(Feb. 8–14, 2016)	423,976 (8.49%)	326,808 (11.66%)	4,260 (0.43%)
	Week 3	5,704,997	3,284,436	1,215,948
	(Feb. 15–21, 2016)	474,652 (8.32%)	373,496 (11.37%)	4,314 (0.35%)
	Week 4	8,580,214	5,071,862	1,661,688
	(Feb. 22–28, 2016)	573,950 (6.69%)	442,897 (8.73%)	4,720 (0.28%)
	Cumulative	23,331,739	13,438,855	3,313,230
		1,788,118 (7.66%)	1,379,016 (10.26%)	6,776 (0.20%)
Ukraine Conflict Bot Source (% of total)	Week 1	306,544	155,151	75,653
	(Aug. 1–7, 2016)	12,605 (4.11%)	8,059 (5.19%)	1,445 (1.91%)
	Week 2	305,796	141,193	107,281
	(Aug. 8–14, 2016)	13,764 (4.50%)	8,272 (5.86%)	1,200 (1.12%)
	Week 3	381,146	210,047	143,684
	(Aug. 15–21, 2016)	17,012 (4.46%)	11,212 (5.34%)	1,647 (1.15%)
	Week 4	280,761	133,985	122,845
	(Aug. 22–28, 2016)	8,772 (3.12%)	5,247 (3.92%)	1,126 (0.92%)
Cumulative	1,274,247	640,376	364,422	
	52,153 (4.09%)	32,790 (5.12%)	2,436 (0.67%)	
Turkey Censorship Bot Source (% of total)	Week 1	709,530	442,429	305,239
	(Dec. 1–7, 2016)	18,243 (2.57%)	11,976 (2.71%)	1,892 (0.62%)
	Week 2	894,900	591,209	410,035
	(Dec. 8–14, 2016)	27,349 (3.06%)	18,189 (3.08%)	2,705 (0.66%)
	Week 3	1,486,289	1,036,436	635,535
	(Dec. 15–21, 2016)	44,833 (3.02%)	29,696 (2.87%)	3,807 (0.60%)
	Week 4	917,937	571,393	425,435
	(Dec. 22–28, 2016)	25,575 (2.79%)	17,379 (3.04%)	2,702 (0.64%)
Cumulative	4,008,656	2,641,467	1,322,010	
	116,000 (2.89%)	77,240 (2.92%)	5,174 (0.39%)	

four-week OSN conversation of interest produces a fairly large cumulative directed retweet network to analyze: 2,431,030 nodes / 8,437,925 edges (U.S. Election), 238,714 nodes / 509,614 edges (Ukraine Conflict) and 1,030,381 nodes / 2,088,524 edges (Turkish Censorship).

Data analysis

We conclude the introduction of our social bot analysis methodological framework by discussing the last step, data analysis. While data analysis is an entirely broad characterization of a step, it is the noted culmination point of acquisition, normalization, fusion and transformation of the harvested Twitter conversation data that enables us to address the overall research questions by applying the methods put forth in the

subsequent sections comprising the “[Analysis results and discussion](#)” of this study. Furthermore, as we seek to contribute to the expansion of social bot analysis techniques, we do not portend that the analysis methods we propose are comprehensive, but merely foundational building blocks paving the way for future application methods.

Analysis results and discussion

In this section, we present the findings of the comparative descriptive statistical analysis methods and social network analysis techniques of this study and discuss the resulting implications of social bot evidence across the global event conversations of interest. By analyzing multiple significant global OSN conversations we expand current social bot analysis literature. Further, we show how the adaptation of SNA techniques can provide quantifiable and comparative results to determine the relative impact or influence of suspected social bots in OSN conversations. “[Bot and human user communication participation](#)” compares the communication trends of humans and bot Twitter users by observing participation volume and identifying the proclivity to engage with certain types of users. “[Temporal persistence of bot centrality rankings](#)” conducts centrality measurements within the retweet networks and evaluates the persistence of social centrality rankings over time. This section concludes with “[Prominent bot ego networks](#)” dissecting the associated ego networks of the highest ranking eigenvector centrality bot from each OSN conversation.

Bot and human user communication participation

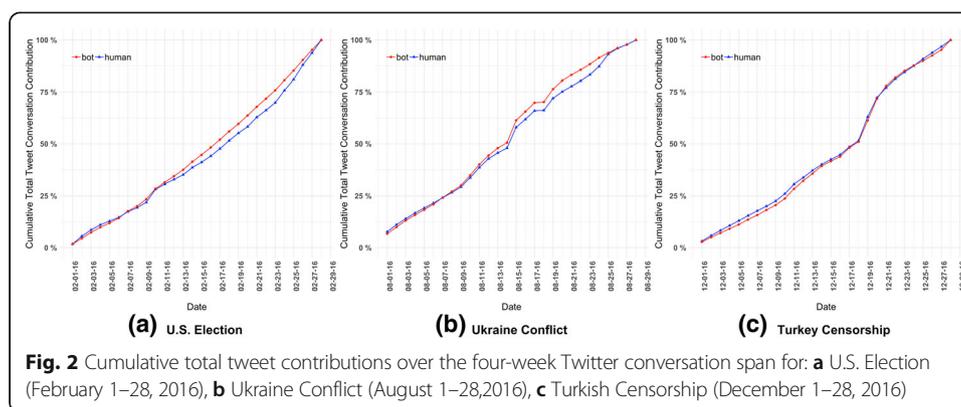
We first compare the communication participation patterns of bot and human users by examining the associated tweet and retweet volume rates. Table 3 summarizes the corresponding average and median volume rates across all three OSNs. We see social bots exhibit much higher average and median participation rates, which is not surprising given the large volume of contributions made by such a small bot population. Of interest though, we see that the top human user account tweet volumes dominate the top bot account tweet volumes across all OSNs, while top bot account retweet volumes are dominant except in the case of the Turkish Censorship OSN.

Figure 2 presents the cumulative total tweet contribution percentages by human and bot users over the four-weeks of harvested tweets for each OSN conversation. The U.S. Election (Fig. 2a) and the Ukraine Conflict (Fig. 2b) conversations both exhibit a gap between bot and human contribution percentages that begins to widen at approximately 2 weeks into the conversation and closes over the final days. A similar gap between users does not exist in the Turkish Censorship conversation (Fig. 2c), while its initial conversation trajectory is much shallower until a spike in contributions takes place corresponding to the onset of the first censorship event in Turkey on December 19, 2016. This latter contribution spike, coupled with lower overall social bot tweet/retweet volumes and participation rates (Tables 2 and 3), might be symptomatic of the Turkish Censorship conversation being an emergent topic during the period of observations, as opposed to the already established U.S. Election and Ukraine Conflict conversations.

The volume of in-group and cross-group communication within OSN retweet conversations provides an additional opportunity to classify communication patterns. We define in-group communication as retweets between like types of users (i.e. humans retweeting humans and bots retweeting bots), while cross-group communication

Table 3 Overall conversation tweet contribution volumes by human and likely social bot users within each OSN conversation corpus of interest

Corpus	User Type	Average User Tweet Volume	Median User Tweet Volume	Min/Max User Tweet Volume (# of users)	Average User Retweet Volume	Median User Retweet Volume	Min/Max User Retweet Volume (# of users)
U.S. Election	Humans	6.5	1.0	Min: 1 (1,920,647 users) Max: 45,542 (1 user)	5.4	1.0	Min: 1 (1,428,380 users) Max: 8,016 (1 user)
	Bots	263.9	13.0	Min: 1 (1,032 users) Max: 19,905 (1 user)	267.0	10.0	Min: 1 (950 users) Max: 19,905 (1 user)
Ukraine Conflict	Humans	3.4	1.0	Min: 1 (231,623 users) Max: 28,072 (1 user)	2.8	1.0	Min: 1 (146,879 users) Max: 1,265 (1 user)
	Bots	21.4	4.0	Min: 1 (569 users) Max: 2,623 (1 user)	168	3.0	Min: 1 (523 users) Max: 2,508 (1 user)
Turkish Censorship	Humans	2.9	1.0	Min: 1 (882,172 users) Max: 10,065 (1 user)	2.7	1.0	Min: 1 (647,261 users) Max: 28,072 (1 user)
	Bots	22.4	4.0	Min: 1 (1,416 users) Max: 1,936 (1 user)	17.3	4.0	Min: 1 (1,331 users) Max: 4,966 (1 user)

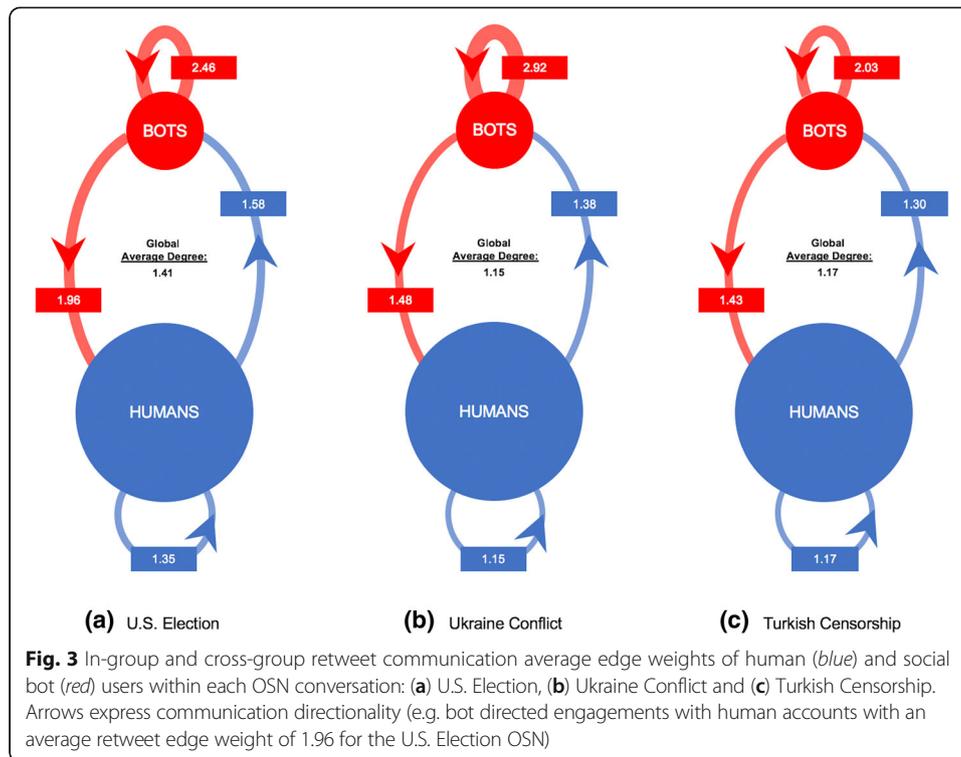


denotes retweets between different types of users (i.e. humans retweeting bots or bots retweeting humans). In terms of total retweet volume percentage for each conversation, humans dominantly retweet other human accounts at total volume rates of 84.92% (U.S. Election), 92.12% (Ukraine Conflict) and 94.74% in (Turkish Censorship), while bot in-group retweet rates occur at relatively low rates of 1.38% and lower. To overcome the human dominance volumes, we normalized retweet interactions by average edge weight of specified group pairings. Figure 3 summarizes the resulting average weighted edges of all inter-group and cross-group communication pairs for each of the OSN conversations of interest. We see that bots, from an average edge weight perspective, engage in higher intra-group and cross-group communication rates across all three conversations, with the U.S. Election conversation showing the highest cross-group and intra-group engagement edge weights of 1.96 and 2.46, respectively.

To further place these overall in-group and cross-group interactions into context, we present the average retweet edge weight for all communication pairings over time in Table 4. The results show that from the weekly and cumulative perspective social bots engage with their in-group bot and cross-group human edge pairs at higher rates, except for the third week of the Turkish Censorship conversation. These across the board higher rates suggest social bots, on average, are hyper-social in comparison to humans: they are much more persistent in attempting to initiate contact with other users in Twitter as opposed to average human users.

Temporal persistence of bot centrality rankings

SNA centrality measurements allow for the derivation of relative node importance, or prominence, based on given node's position in the structure of the network relative to other nodes (Wasserman and Faust 1994). This section employs centrality measurements to determine the relative importance of social bots within each of this study's OSN conversations. As Riquelme and González-Cantergiani (2016) discusses, there exist a vast quantity of available centrality measurements to measure the influence of a particular user in a directed retweet network of a particular Twitter conversation. In this study, we purposely selected the following six centrality measurements due to their relatively common recognition and efficient computational requirements: (1) degree, (2) in-degree, (3) out-degree, (4) eigenvector, (5) betweenness and (6) PageRank.



Degree centrality is the total number of direct edges a node shares with other nodes in a network and does not recognize edge directionality. In a retweet network, degree centrality is synonymous with a Twitter user's popularity in the network. In-degree and out-degree centrality are simply degree centrality that take into account edge directionality. Nodes with higher in-degree centrality receive more directional edge

Table 4 Average retweet edge weight for all inter-group and cross-group communications by human and bot users for each OSN conversation

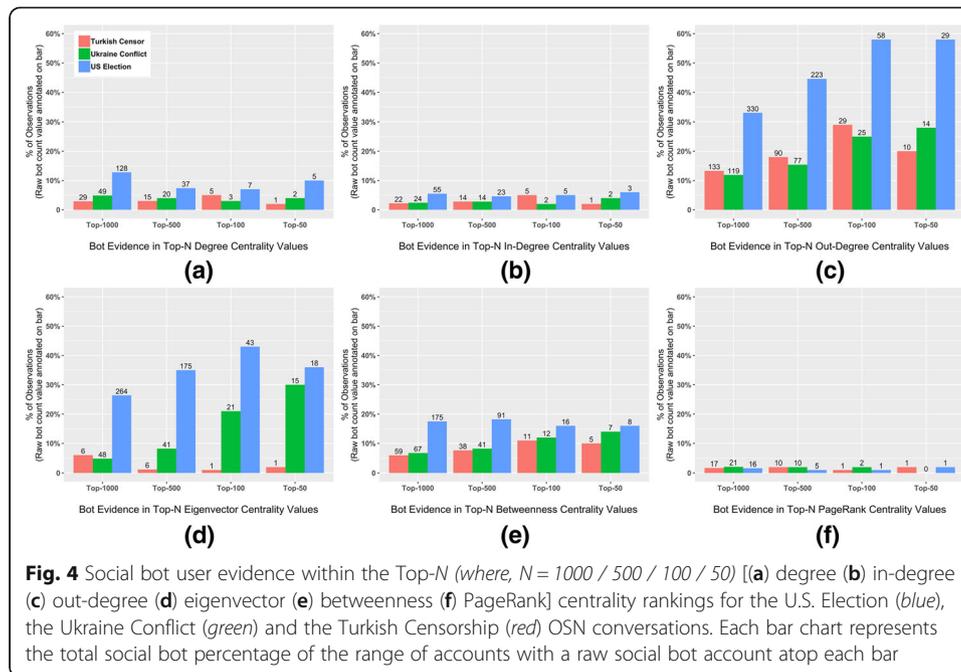
Corpus	Week(s)	Average Retweet Edge Weight			
		Bot-to-Bot	Bot-to-Human	Human-to-Bot	Human-to-Human
U.S. Election	Week 1 (Feb. 1–7, 2018)	2.04	1.83	1.44	1.29
	Week 2 (Feb. 8–14, 2018)	2.61	2.08	1.69	1.47
	Week 3 (Feb. 15–21, 2018)	2.68	1.94	1.61	1.32
	Week 4 (Feb. 22–28, 2018)	2.42	1.96	1.58	1.35
	<i>Cumulative</i>	2.46	1.96	1.58	1.35
Ukraine Conflict	Week 1 (Aug. 1–7, 2018)	3.02	1.47	1.39	1.15
	Week 2 (Aug. 8–14, 2018)	3.84	1.71	1.55	1.21
	Week 3 (Aug. 15–21, 2018)	2.42	1.32	1.31	1.12
	Week 4 (Aug. 22–28, 2018)	2.64	1.58	1.36	1.15
	<i>Cumulative</i>	2.92	1.48	1.38	1.15
Turkish Censorship	Week 1 (Dec. 1–7, 2018)	2.78	1.46	1.28	1.21
	Week 2 (Dec. 8–14, 2018)	1.85	1.47	1.26	1.16
	Week 3 (Dec. 15–21, 2018)	1.88	1.38	1.25	1.14
	Week 4 (Dec. 22–28, 2018)	2.23	1.47	1.48	1.19
	<i>Cumulative</i>	2.03	1.43	1.30	1.17

contact from other nodes, while higher out-degree centrality signifies nodes that initiate more directional edge contact. In a retweet network, higher out-degree centrality equates to a Twitter user initiating more retweets, while higher in-degree means a Twitter user has more users retweeting its original messages. Eigenvector centrality is the weighted sum of all direct and indirect edges for a node that takes into account the individual degree centrality of each node in the network (Bonacich 2007). From a retweet network perspective, eigenvector centrality is a global measure of influence within a conversation. Betweenness centrality measures the propensity of a given node falling on the shortest path between all other node pairs in a network (Freeman 1977). We can view the betweenness centrality of a retweet network node as a measure of communication that flows through that specific node. Finally, PageRank is a derivation of eigenvector centrality, but places more importance on the degree value of the nodes that initiate edges with a node of interest (Brin and Page 1998). Therefore, in a retweet network, a node with higher PageRank value receives more retweets from Twitter users that have greater popularity in the network.

To determine the relative importance of social bot users compared to human users, we calculate the chosen centrality measures for the entire duration of each OSN conversation using the applicable centrality functions provided in the *networkx* Python package (Hagberg et al. 2008). Scale tests by the authors on larger Twitter datasets of at least twice the volumes of the events in this study (i.e. ~ 50 million tweets) comprising networks with cumulative edge volumes that are three times larger (i.e. ~ 25 million edges) returned efficient centrality processing times (i.e. PageRank calculation was most time intensive calculation at ~ 5 min 20 s) within in a cloud environment with the same specifications detailed in the ‘Data acquisition and processing’ section. We then rank order and present the density of social bots within the Top- N centrality ranking positions (where, $N = 1000 / 500 / 100 / 50$). The results (Fig. 4) clearly show that suspected bot users, while representing only 0.28% of all corpus users, account for a significant number of high centrality rankings, especially out-degree and eigenvector centrality rankings. The prevalence of social bots among the top ranks of out-degree nodes shows the above-mentioned hyper-social attitude of bots: they attempt to induce interaction by retweeting other users at a significantly higher rate than their human counterparts. In terms of influence, we see bots infiltrate some of the highest eigenvector centrality rankings within the U.S. Election and the Ukraine Conflict conversations, where bots account for 36.0% and 30.0% of the Top-50 influential accounts, respectively. These results are quite substantial given the employment of just one bot detection source.

To evaluate the temporal persistence of social bot centrality rankings, we recalculate and directly compare centrality rankings in a cumulative fashion over the 4 weeks for each OSN conversation. In doing so, we are able to analyze the centrality ranking staying power of identified social bot accounts over time, as opposed to an overall snapshot of the entire corpus timeframe. Figure 5 (U.S. Election), Fig. 6 (Ukraine Conflict) and Fig. 7 (Turkish Censorship) present a consolidated visualization depicting the density of bot (*red block*) and human (*blue block*) users as each conversation progresses on a weekly cumulative basis, while also annotating the individual accounts within each block.

The centrality ranking persistence of suspected bot users is visually evident over time across the cumulative conversations. We see persistent bot density within each centrality ranking with especially high density associated with the out-degree and eigenvector



centralities for the U.S. Election and the Ukraine Conflict conversations. This includes social bots achieving extremely high-rankings to include two of the top-5 out-degree, eigenvector and centrality rankings within the Ukraine Conflict conversation (Fig. 5) and seven and four of the top-10 out-degree and eigenvector centrality rankings, respectively, within the U.S. Election conversation (Fig. 6).

Observing the classification results of popular news source accounts (e.g. @CNN, @thehill, @AP) highlights the shortcomings of using only one bot detection service. For example, DeBot classifies @thehill as an automated bot account, but does not for @AP or @CNN. We can only assume, therefore, that coverage by DeBot has not evaluated those accounts by the time of this study. The account @FoxNews was later evaluated after this study by DeBot and determined to be an automated account on May 5, 2018, but we maintained its original label given the evaluation dates of this study. Further extensions of this proof-of-concept work should include additional bot detection services, while consideration should be taken into potentially removing verified accounts from evaluation.

Prominent bot ego networks

In this final “[Analysis results and discussion](#)” section, we investigate the ego networks of the highest ranking eigenvector centrality social bots from the U.S. Election (Twitter ID: 732980827, Username: *ChristiChat*) and Ukraine Conflict (Twitter ID: 3346642625, Username: *justfightX*) OSN retweet conversations. No Turkish social bots achieved a high sustained eigenvector centrality ranking, so we did not include a Turkish bot in this section. Using the *ego_graph* function provided in the Python *networkx* package, we derived the ego-networks based on immediately adjacent neighbors for each of the identified accounts. Formally, we evaluated the individual sub-graph network comprised of only immediate (i.e.

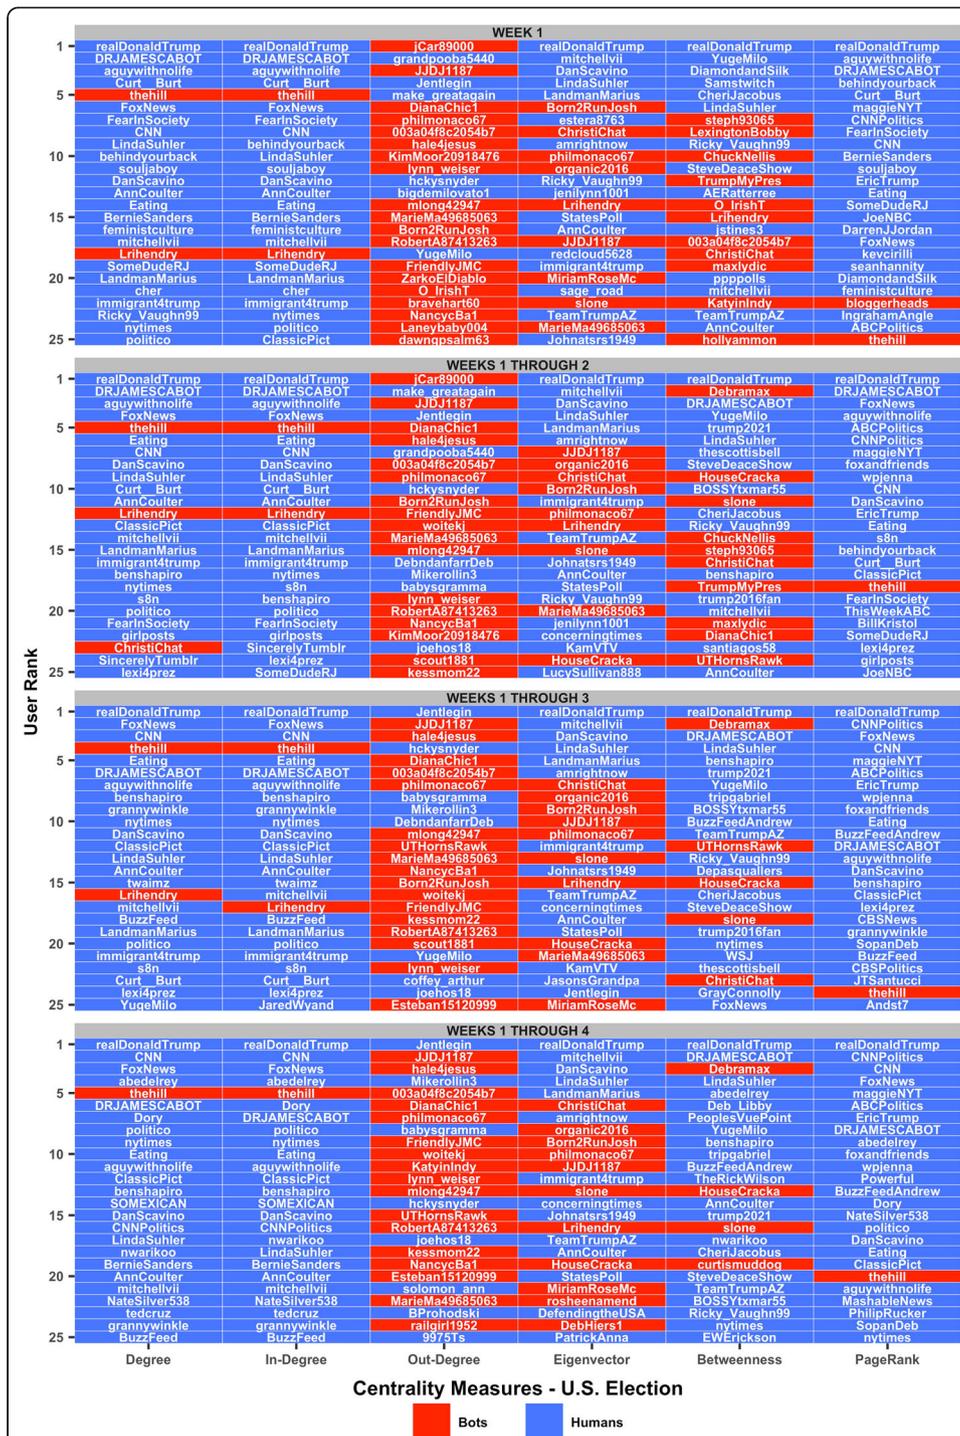
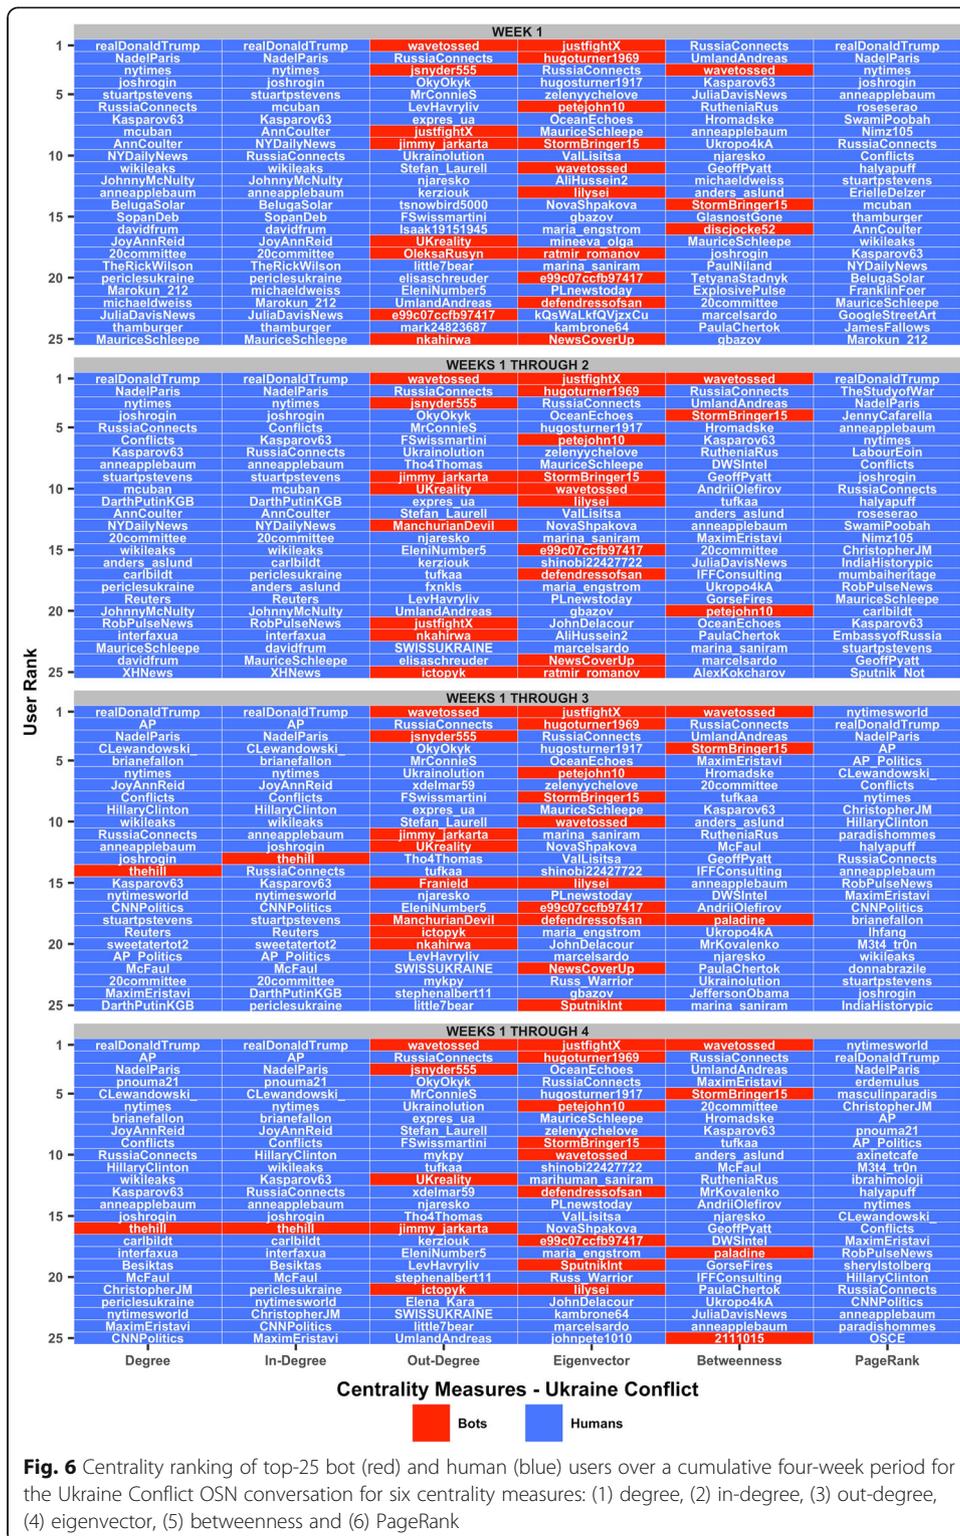


Fig. 5 Centrality ranking of top-25 bot (red) and human (blue) users over a cumulative four-week period for the U.S. Election OSN conversation for six centrality measures: (1) degree, (2) in-degree, (3) out-degree, (4) eigenvector, (5) betweenness and (6) PageRank

directly connected) neighbors for the highest ranking eigenvector centrality social bot accounts within the larger specific OSN graph network consisting of all accounts. We extract the observable retweet network characteristics of these most relatively influential social bot users and directly compare them to the average retweet bot characteristics



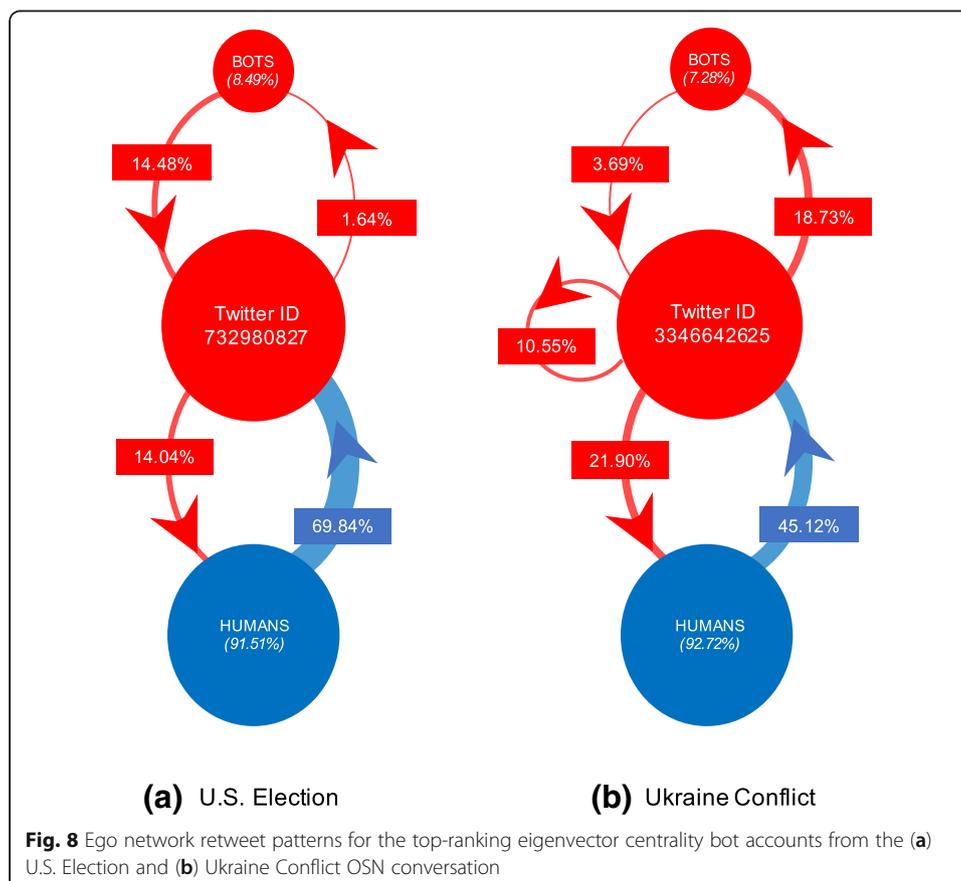
presented in the “Bot and human user communication participation” section. Figure 8 provides a proportionally-scaled ego network that depicts the inter-group and cross-group neighbor interactions of these top eigenvector social bots. While both of these influential bots engage in differing levels of inter-group communication with other bots and

humans accounting for retweet rates 69.84% and 45.12% within the U.S. Election and Ukraine Conflict ego networks, respectively.

Conclusion

In this study, we present a novel approach to expand the emergent area of social bot research. The unique social bot analysis methodological framework put forth enables the inclusion of additional bot detection platform services, while also opening the analysis window to account for new OSN conversations of interest. Through the lens of three major global event OSN conversations in 2016, we confirmed the hyper-social nature of bots: suspected social bots users make far more attempts on average than human users to initiate contact with other users via retweets. Social network analysis centrality measurements discover that social bots, while comprising less than 0.3% of the total user population, display a profound level of structural network influence by ranking particularly high among the top eigenvector centrality users within the U.S. Election and the Ukraine Conflict OSN conversations. Further, we determine that social bots exhibit temporal persistence in centrality ranking density across all of the OSN conversation.

While we report promising findings, this study must account for its many limitations. Relying upon a single bot detection platform helped validate this study's applied network analysis methods, but a sole source detection algorithm is not sufficient for overcoming known specific limitations that currently challenge all open-source bot detection results (Subrahmanian et al. 2016; Cresci et al. 2017). Also, solely using data from a single OSN



platform induces a litany of associated biases to include representativeness and sampling bias shortcomings (Tufekci 2014). Ruths and Pfeffer (2014) further expands on social media data issues, while also singling out the inability to properly determine the presence of bots. While it is also in the spirit of this study to help improve overall bot detection methods, it is a reasonable perspective to state the current difficulties to determine ground truth effectiveness in detecting bots (Subrahmanian et al. 2016; Cresci et al. 2017; Chavoshi and Mueen 2018). Further, a binary classification between bots and humans is not entirely sufficient as cyborg accounts also exist, which Chu et al. (2012) coins as bot-assisted human or human-assisted bot account.

Immediate primary extensions of this work should expand beyond the proof concept framework demonstrated here and aggressively seek the inclusion of additional bot detection algorithms for a more holistic bot labeling perspective. While there currently exists a limited number of open-source bot detection algorithms, a comprehensive collection of detection sources would ideally include access to the continually improving pre-existing detection platforms (Varol et al. 2017; Chavoshi et al. 2017; Beskow and Carley 2018), as well as recent novel detection algorithms based on detecting evolving bot signatures (Cresci et al. 2018c; Mazza et al. 2019). Further extensions of this work could aim to incorporate additional social media sources beyond Twitter as Hecking et al. (2018) describe in a cross-media information diffusion example sourcing data from Twitter, Wikipedia edits and other web-based sources. In the case of this study, if we do not observe centrality measures beyond just degree and PageRank centrality, then we miss the important social rankings made available via out-degree and eigenvector centrality. Therefore, it is important to maintain an expansive centrality analysis to account for social bots by potentially incorporating additional centrality measures, such as percolation centrality (Piraveenan et al. 2013), that may perform well in ranking social bot prominence within networks. On its own, this paper is a unique stepping stone that adds to the growing research efforts focused on understanding social bot behavior in global event conversations.

Abbreviations

API: Application programming interface; AWS: Amazon web services; CSS: Computational social science; HITS: Hypertext induced topic search; LTRSB: Learning to rank social bots; OSN: Online social network; SNA: Social network analysis; TC: Turkish censorship; UC: Ukraine conflict; UE: United States election; U.S: United States

Acknowledgements

The authors would like to thank the DeBot team, led by Nikan Chavoshi, at the University of New Mexico for providing complete access to the DeBot platform bot archive.

Authors' contributions

RS, ATC and AC developed the topic, and RS along with ATC and AS prepared the initial draft; RS, ATC, AS and AC prepared the final draft and all authors approved its content.

Funding

The authors received no funding for this research.

Availability of data and materials

The datasets generated and analyzed during the current study are available in the UNC Dataverse repository, which is accessible at <https://doi.org/10.15139/S3/SGZQGT>.

Competing interests

The authors declare that they have no competing interests.

Author details

¹Computational Social Science Program, Department of Computational and Data Sciences, George Mason University, Fairfax, Virginia 22030, USA. ²Department of Geography and Geoinformation Science, George Mason University, Fairfax,

Virginia 22030, USA. ³Criminal Investigations and Network Analysis Center, George Mason University, Fairfax, Virginia 22030, USA.

Received: 28 February 2019 Accepted: 1 July 2019

Published online: 01 August 2019

References

- Abokhodair N, Yoo D, McDonald DW (2015) Dissecting a social botnet: growth, content and influence in twitter. In: Proc. of 18th ACM CSCW 2015, pp 839–851
- Aiello LM, Deplano M, Schifanella R, Ruffo G (2014) People are strange when you're a stranger: impact and influence of bots on social networks. In: Proc. of the 6th AAAI international Conf. On weblogs and social media. AAAI, Dublin, pp 10–17
- Avvenuti M, Cresci S, Marchetti A et al (2016a) Predictability or early warning: using social media in modern emergency response. *IEEE Internet Comput* 20:4–6
- Avvenuti M, Cresci S, Vigna FD, Tesconi M (2016b) Impromptu crisis mapping to prioritize emergency response. *Computer* 49:28–37
- Bakshy E, Hofman JM, Mason WA, Watts DJ (2011) Everyone's an influencer: quantifying influence on twitter. In: Proc of the Fourth ACM International Conf on Web Search and Data Mining. ACM, New York, pp 65–74
- Beskow DM, Carley KM (2018, 2018) Bot-hunter: a tiered approach to Detecting & Characterizing Automated Activity on twitter. SBP-BRIMS 2018. Intl. Conf. on Social Computing, Behavioral-Cultural Modeling and Prediction and Behavior Representation in Modeling and Simulation
- Bessi A, Ferrara E (2016) Social bots distort the 2016 U.S. presidential election online discussion. *First Monday* 21(11)
- Blackwell D, Leaman C, Tramosch R et al (2017) Extraversion, neuroticism, attachment style and fear of missing out as predictors of social media use and addiction. *Personal Individ Differ* 116:69–72
- Bonacich P (2007) Some unique properties of eigenvector centrality. *Soc Networks* 29:555–564
- Boshmaf Y, Muslukhov I, Beznosov K, Ripeanu M (2013) Design and analysis of a social botnet. *Comput Netw* 57:556–578
- Boyd D, Golder S, Lotan G (2010) Tweet, tweet, retweet: conversational aspects of retweeting on twitter. In: Proceedings of the 2010 43rd Hawaii international conference on system sciences. IEEE Computer Society, Washington, pp 1–10
- Brin S, Page L (1998) The anatomy of a large-scale hypertextual web search engine. *Comput Netw ISDN Syst* 30:107–117
- Broniatowski DA, Jamison AM, Qi S et al (2018) Weaponized health communication: twitter bots and Russian trolls amplify the vaccine debate. *Am J Public Health* 108:1378–1384
- Cha M, Haddadi H, Benevenuto F, Gummadi KP (2010) Measuring user influence in twitter : the million follower fallacy. In: Proceedings of the fourth international AAAI conference on weblogs and social media (ICWSM 2010). AAAI Press, Washington, DC, pp 10–17
- Chavoshi N, Hamooni H, Mueen A (2016) DeBot: twitter bot detection via warped correlation. In: 2016 IEEE 16th International Conference on Data Mining (ICDM), pp 817–822
- Chavoshi N, Hamooni H, Mueen A (2017) Temporal patterns in bot activities. In: Proc. of 26th International Conf. on WWW, pp 1601–1606
- Chavoshi N, Mueen A (2018) Model bots, not humans on social media. In: 2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM) IEEE, pp 178–185
- Chu Z, Gianvecchio S, Wang H, Jajodia S (2012) Detecting automation of twitter accounts: are you a human, bot, or cyborg? *IEEE Trans Dependable Secure Comput* 9:811–824
- Ciampaglia GL (2018) Fighting fake news: a role for computational social science in the fight against digital misinformation. *J Comput Soc Sc* 1:147–153
- Conover MD, Ratkiewicz J, Francisco M et al (2011) Political polarization on twitter. In: Fifth international AAAI conference on weblogs and social media, pp 10–17
- Cresci S, Di Pietro R, Petrocchi M et al (2017) 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, Republic and Canton of Geneva, pp 963–972
- Cresci S, Lillo F, Regoli D et al (2018a) \$ FAKE: evidence of spam and bot activity in stock microblogs on twitter. In: Twelfth international AAAI conference on web and social media, pp 580–583
- Cresci S, Lillo F, Regoli D et al (2019a) Cashtag piggybacking: uncovering spam and bot activity in stock microblogs on twitter. *ACM Trans Web* 13:11:1–11:27
- Cresci S, Petrocchi M, Spognardi A, Tognazzi S (2018b) From reaction to Proaction: unexplored ways to the detection of evolving Spambots. In: WWW (Companion Volume), pp 1469–1470
- Cresci S, Petrocchi M, Spognardi A, Tognazzi S (2019b) On the capability of evolved spambots to evade detection via genetic engineering. *Online Soc Netw Media* 9:1–16
- Cresci S, Petrocchi M, Spognardi A, Tognazzi S (2019c) Better safe than sorry: an adversarial approach to improve social bot detection. In *WebSci '19: Proc. of the 11th ACM Conference on Web Scienc.* ACM, New York, pp 47–56
- Cresci S, Pietro RD, Petrocchi M et al (2018c) Social fingerprinting: detection of Spambot groups through DNA-inspired behavioral modeling. *IEEE Trans Dependable Secure Comp* 15:561–576
- Crooks A, Croitoru A, Stefanidis A, Radzikowski J (2013) #earthquake: twitter as a distributed sensor system. *Trans GIS* 17:124–147
- Davis CA, Varol O, Ferrara E et al (2016) In WWW '16 Companion: Proc. of the 25th Intl. Conf. Companion on World Wide Web, IW3C2, Geneva, pp 273–274
- Duh A, Slak Rupnik M, Korošak D (2018) Collective behavior of social bots is encoded in their temporal twitter activity. *Big Data* 6:113–123
- Ferrara E (2017) Contagion dynamics of extremist propaganda in social networks. *Inf Sci* 4(18):1–12
- Ferrara E, Varol O, Davis C et al (2016) The rise of social bots. *Commun ACM* 59:96–104
- Freeman LC (1977) A set of measures of centrality based on Betweenness. *Sociometry* 40:35–41
- Fuchs C (2005) The internet as a self-organizing socio-technological system. *Cybernetics Human Knowing* 12:37–81
- Grinberg N, Joseph K, Friedland L et al (2019) Fake news on twitter during the 2016 U.S. presidential election. *Science* 363:374–378

- Hagberg A, Schult D, Swart P (2008) Exploring network structure, dynamics, and function using NetworkX. In: SciPy2008: Proc. of the 7th Python in science conference, pp 11–15
- Hecking T, Steiner L, Masias VH, Ulrich Hoppe H (2018) Relational patterns in cross-media information diffusion networks. In: Cherifi C, Cherifi H, Karsai M, Musolesi M (eds) *Complex Networks & Their Applications VI*. Springer International Publishing, Cham, pp 1002–1014
- Hegelich S, Janetzko D (2016) Are Social Bots on Twitter Political Actors? Empirical Evidence from a Ukrainian Social Botnet. In: Proc. Of the 10th Intl. Conf. on Web and Social Media (ICWSM), ICWSM, pp 579–582
- Howard PN, Kollanyi B (2016) Bots, #StrongerIn, and #Brexit: computational propaganda during the UK-EU referendum. SSRN, <https://doi.org/10.2139/ssrn.2798311>
- Howard PN, Woolley S, Calo R (2018) Algorithms, bots, and political communication in the US 2016 election: the challenge of automated political communication for election law and administration. *J Inform Tech Polit* 15:81–93
- Kwak H, Lee C, Park H, Moon S (2010) What is twitter, a social network or a news media? In: Proceedings of the 19th international conference on world wide web. ACM, New York, pp 591–600
- Lazer DMJ, Baum MA, Benkler Y et al (2018) The science of fake news. *Science* 359:1094–1096
- Mazza M, Cresci S, Avenuti M et al (2019) RTbust: exploiting temporal patterns for botnet detection on twitter. In *WebSci '19: Proc. of the 11th ACM Conference on Web Science*. ACM, New York, pp 183–192
- Mitchell A (2018) Americans still prefer watching to Reading the news - and mostly still through television. Pew Research Center, Washington, D.C.
- Mønsted B, Sapiezzyński P, Ferrara E, Lehmann S (2017) Evidence of complex contagion of information in social media: an experiment using twitter bots. *PLoS One* 12:e0184148
- Morstatter F, Wu L, Nazer TH et al (2016) A new approach to bot detection: striking the balance between precision and recall. In: 2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM), pp 533–540
- Murthy D, Powell AB, Tinati R et al (2016) Automation, algorithms, and politics| bots and political influence: a sociotechnical investigation of social network capital. *Int J Commun* 10:20
- Perna D, Tagarelli A (2018) Learning to rank social bots. In: Proceedings of the 29th on hypertext and social media. ACM, New York, pp 183–191
- Piraveenan M, Prokopenko M, Hossain L (2013) Percolation centrality: quantifying graph-theoretic impact of nodes during percolation in networks. *PLoS One* 8(1):e53095
- Reece AG, Reagan AJ, Lix KLM et al (2017) Forecasting the onset and course of mental illness with twitter data. *Sci Rep* 7:13006
- Riquelme F, González-Cantergiani P (2016) Measuring user influence on twitter: a survey. *Inf Process Manag* 52:949–975
- Ruths D, Pfeffer J (2014) Social media for large studies of behavior. *Science* 346:1063–1064
- Sakaki T, Okazaki M, Matsuo Y (2013) Tweet analysis for real-time event detection and earthquake reporting system development. *IEEE Trans Knowl Data Eng* 25:919–931
- Schuchard R, Crooks A, Stefanidis A, Croitoru A (2019) Bots in nets: empirical comparative analysis of bot evidence in social networks. In: Aiello LM, Cherifi C, Cherifi H et al (eds) *Complex networks and their applications VII*. Springer International Publishing, Cham, pp 424–436
- Shao C, Ciampaglia GL, Varol O et al (2018) The spread of low-credibility content by social bots. *Nat Commun* 9:4787
- Stella M, Ferrara E, Domenico MD (2018) Bots increase exposure to negative and inflammatory content in online social systems. *PNAS* 115:12435–12440
- Strohmaier M, Wagner C (2014) Computational social science for the world wide web. *IEEE Intell Syst* 29:84–88
- Suárez-Serrato P, Roberts ME, Davis C, Menczer F (2016) On the influence of social bots in online protests. In: Spiro E, Ahn Y-Y (eds) *Social informatics*. Springer International Publishing, Berlin, pp 269–278
- Subrahmanian VS, Azaria A, Durst S et al (2016) The DARPA twitter bot challenge. *Computer* 49:38–46
- Sunstein CR (2018) *#republic: divided democracy in the age of social media*. Princeton University Press, Princeton, NJ
- Tufekci Z (2014) Big questions for social media big data: representativeness, validity and other methodological pitfalls. In: ICWSM '14: Proc. of the 8th Intl. AAAI Conference on Weblogs and Social Media. AAAI, Palo Alto, pp 505–514.
- Varol O, Ferrara E, Davis CA et al (2017) Online human-bot interactions: detection, estimation, and characterization. In: Proc. of the 11th international AAAI Conf. On web and social media. AAAI, Montréal, pp 280–289
- Vosoughi S, Roy D, Aral S (2018) The spread of true and false news online. *Science* 359:1146–1151
- Wasserman S, Faust K (1994) *Social network analysis: methods and applications*, 1st edn. Cambridge University Press, Cambridge
- Weng J, Lim E-P, Jiang J, He Q (2010) TwitterRank: finding topic-sensitive influential Twitterers. In: Proceedings of the Third ACM International Conference on Web Search and Data Mining. ACM, New York, pp 261–270

Publisher's Note

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