



ARTICLE

<https://doi.org/10.1057/s41599-019-0359-x>

OPEN

Bots fired: examining social bot evidence in online mass shooting conversations

Ross Schuchard^{1*}, Andrew Crooks^{1,2}, Anthony Stefanidis³ & Arie Croitoru²

ABSTRACT Mass shootings, like other extreme events, have long garnered public curiosity and, in turn, significant media coverage. The media framing, or topic focus, of mass shooting events typically evolves over time from details of the actual shooting to discussions of potential policy changes (e.g., gun control, mental health). Such media coverage has been historically provided through traditional media sources such as print, television, and radio, but the advent of online social networks (OSNs) has introduced a new platform for accessing, producing, and distributing information about such extreme events. The ease and convenience of OSN usage for information within society's larger growing reliance upon digital technologies introduces potential unforeseen risks. Social bots, or automated software agents, are one such risk, as they can serve to amplify or distort potential narratives associated with extreme events such as mass shootings. In this paper, we seek to determine the prevalence and relative importance of social bots participating in OSN conversations following mass shooting events using an ensemble of quantitative techniques. Specifically, we examine a corpus of more than 46 million tweets produced by 11.7 million unique Twitter accounts within OSN conversations discussing four major mass shooting events: the 2017 Las Vegas concert shooting, the 2017 Sutherland Springs church shooting, the 2018 Parkland School Shooting and the 2018 Santa Fe school shooting. This study's results show that social bots participate in and contribute to online mass shooting conversations in a manner that is distinguishable from human contributions. Furthermore, while social bots accounted for fewer than 1% of total corpus user contributors, social network analysis centrality measures identified many bots with significant prominence in the conversation networks, densely occupying many of the highest eigenvector and out-degree centrality measure rankings, to include 82% of the top-100 eigenvector values of the Las Vegas retweet network.

¹Computational Social Science Program, Department of Computational and Data Sciences, George Mason University, Fairfax, USA. ²Department of Geography and Geoinformation Science, George Mason University, Fairfax, VA, USA. ³Criminal Investigations and Network Analysis Center, George Mason University, Fairfax, VA, USA. *email: rschucha@gmu.edu

Introduction

Mass shootings have become their own distinct phenomenon separate from the likes of general homicide and mass murder due to their continued prevalence and the natural draw of media attention to extreme events (Schildkraut et al., 2018). While not a formally defined government statistic, a mass shooting has generally been defined as an incident resulting in the death of four or more victims, not including the killer (Towers et al., 2015; Dahmen et al., 2018; Silva and Capellan, 2019). Moffat (2019) tallies that 620 people have been killed and more than 1,000 wounded from 70 mass shooting events in the United States beginning with the Columbine shooting in 1999 and concluding with the Parkland shooting in 2018. While the media reporting environment has changed drastically since the Columbine shooting with the advent of online social networks (OSNs) driven by the Web 2.0 paradigm, the general public's interest in mass shooting coverage remains high given the recent historical increase in mass shootings and the associated debate on the polarising topic of gun control (Newman and Hartman, 2017). Media research has shown that particular newsworthy events, such as mass shootings, lend themselves to framing, whereby the purposive highlighting of certain attributes of a single event may attract or sustain specific interest or viewpoints of the general public when it processes that information (e.g., Entman, 1993; Chyi and McCombs, 2004; Chong and Druckman, 2007). Guggenheim et al. (2015) points to a reciprocal relationship in framing mass shooting narratives between traditional and OSN media sources, but also highlights how OSN content creators and users transcend the typical journalistic gatekeeping norms of traditional media by openly expressing emotional reactions.

In the United States, OSNs recently surpassed traditional print newspapers as a primary source for news and continue to gain traction on other traditional news sources such as television and radio (Mitchell, 2018). Mahabir et al. (2018) describes the current news ecosystem, comprised of both traditional news sources and online platforms (e.g., news websites/apps and social media), as highly participatory and fostering digital activism. Edwards et al. (2013) defines digital activism as an organised public effort orchestrated by supporters of an issue using digital media to make collective claims against a target authority. Given the highly contentious policy debates surrounding gun control that are a typical conversational byproduct in the immediate aftermath of a mass shooting event (Merry, 2016; Newman and Hartman, 2017), OSN conversations about mass shootings are a salient topic for digital activists. While a convenient means to access and publish content, OSNs have proven to be complicit in spreading and amplifying manipulated and/or blatantly falsified narratives (Bolsover and Howard, 2017; Starbird, 2017; Lazer et al., 2018; Vosoughi et al., 2018). A primary factor contributing to the skewed narratives in OSNs is the existence of vast populations of social media accounts controlled by social bots (Boshmaf et al., 2013).

Social bots are computer algorithms that automatically produce content and interact with human OSN users (Ferrara et al., 2016). The pervasiveness of social bots has led to numerous research efforts focused on developing novel bot detection methods (Davis et al., 2016; Chavoshi et al., 2016) and examining how bots spread information (Aiello et al., 2014; Mønsted et al., 2017; Shao et al., 2018). It is necessary to assert that the automated behaviour of bots does not imply intentional malice, as many bots serve in benign roles (e.g., news feed aggregator) (Ferrara et al., 2016). Further introductory works have analysed the presence of social bots in various polarising OSN conversations such as elections (Howard and Kollanyi, 2016; Bessi and Ferrara, 2016), conflict (Schuchard et al., 2019) and vaccinations (Subrahmanian et al., 2016; Broniatowski et al., 2018).

While some promising recent works have touched in various ways on the topic of social bots in OSN mass shooting conversations in various forms (Nied et al., 2017; Starbird, 2017; Kitzie et al., 2018), which will be discussed in more detail in the Background section, there is a need for in-depth quantitative social bot analysis. In light of this, this study expands the literature by introducing quantitative ensemble methods to observe social bot activity in large-scale online conversations by examining suspected social bots within Twitter conversations associated with four recent mass shooting events: the Las Vegas concert shooting (October 1, 2017), the Sutherland Springs church shooting (November 5, 2017), the Parkland School Shooting (February 14, 2018) and the Santa Fe school shooting (May 18, 2018). Specifically, we analysed the presence and contribution patterns of social bots in relation to human users in an effort to determine potential cross-conversational norms of bot behaviour within the highly polarised conversation topic of mass shooting events. We further sought to quantify and classify the mentioning rate of previous mass shooting events in conversations about subsequent events to potentially classify certain events with persistent salience. Finally, we employed social network analysis centrality measures to derive the relative structural importance of social bots in relation to other users within each of the OSN conversation networks.

This study's results show that social bots participate and contribute to online mass shooting conversations in a manner that is distinguishable from human contributions. The cumulative conversation contribution rates of bots outpace humans throughout the Sutherland Springs and Santa Fe conversations. In the conversations involving the highly salient Las Vegas and Parkland shootings, human contributions initially outpace social bots, but an inversion takes place within the first week of each conversation as social bots become the leading contributor for the remainder of the conversation. In terms of cross-group communications, human accounts engaged suspected bot accounts at higher rates than bots engaged humans in all of the conversations. Finally, bots, while accounting for fewer than 1% of all corpus users, displayed significant prominence in the conversation networks, densely occupying many of the highest eigenvector and out-degree centrality measure rankings, including 82% of the top-100 eigenvector values of the Las Vegas tweet network.

The remainder of this paper is as follows. First, we present applicable previous works in the 'Background' section. Next, in the 'Data and methods' section, we present a detailed overview of our data acquisition and processing steps, as well as introducing the methods employed in the study. The 'Results and discussion' section presents the findings of our employed methods to answer the study's research questions, followed by the 'Conclusion' section.

Background

A recent study by Dahmen et al. (2018) found that a majority of journalists agreed that the traditional news coverage of mass shooting events has become routine due to perceived formulaic reporting by reporters. However, some mass shooting events garner more initial and sustained attention than others, and traditional media studies have focused much effort on identifying how certain media reports are able to succeed at this (Schildkraut et al., 2018; Silva and Capellan, 2019). One means of achieving both initial and sustained attention is through dynamic frame changing, or emphasising different aspects of a news event over its life span (Chyi and McCombs, 2004). In terms of mass shootings, Muschert and Carr (2006) applied and extended the dynamic framing concept by analysing frame changes over the

course of media reporting on nine school shootings from 1997 to 2001. This resulted in the ability to directly compare the highly salient Columbine school shooting event with eight less salient shooting events (i.e., Pearl, MS; Paducah, KY; Jonesboro, AR; Edinboro, PA; Springfield, OR; Conyers, GA; Santee, CA; El Cajon, CA) using the frame-changing spatial (community, regional, societal) and temporal (past, present, future) categorisations (shown in Fig. 1). Schildkraut and Muschert (2014) extended the nine-school mass shooting research of Muschert and Carr (2006) to include the Sandy Hook Elementary School shooting in an effort to be able to directly compare it to the Columbine shooting, a precedent setting mass shooting event of extremely high salience for media coverage, while ultimately determining that the Sandy Hook shooting reshaped media coverage of mass shooting events.

Although emerging from traditional media research, the concept of media framing serves as a primary analysis method in OSN conversations involving mass shootings as well. For instance, Guggenheim et al. (2015) examined the framing of mass shooting events in both traditional and OSN media coverage and concluded that there is a reciprocal relationship between the two. That is, tweets respond to traditional media reports just as traditional media reports respond to tweets. To account for additional narrative contributors, Merry (2016) extended the concept of mass shooting narrative framing to include framing by interest groups in OSN Twitter conversations (e.g., National Rifle Association, Brady Campaign to Prevent Gun Violence). In other work, Starbird (2017) examined the propagation and shaping of alternative narratives emanating from tweeted URLs related to mass shooting events, exposing how OSN interactions can enable a conspiratorial ecosystem of alternative media.

Turning our attention to social bots, Varol et al. (2017) estimated that social bots account for 9–15% of all Twitter user accounts. This relatively large bot population estimate and the associated unknown implications of human users engaging with

non-humans in Twitter have led to the rapid emergence of bot detection and classification research. While the ever-increasing sophistication of social bots mimicking human behaviour has proven to be quite difficult phenomenon for researchers to keep pace with (Cresci et al., 2017), multiple bot detection research platforms have been launched to aid researchers in the overall detection and analysis of social bots in Twitter. Botometer¹, formerly named BotOrNot, is a widely used open-source bot detection platform that employs a supervised random forest detection algorithm against more than 1,100 extracted unique account features (Davis et al., 2016; Varol et al., 2017). Botometer then returns a classification score on a normalised scale identifying an account as more “human” or “bot-like”. Chavoshi et al. (2016) developed the open-source DeBot² bot detection platform that employs an unsupervised warped correlation bot detection algorithm. Rather than relying on feature extraction, the DeBot platform provides a binary bot classification based solely on the synchronous temporal activities of Twitter accounts.

While there is ample room for growth beyond the promising initial social bot analysis research identified in the ‘Introduction’ section of this paper, substantial exploratory work is still needed to be started in analysing bot activity in OSN conversations involving mass shootings. Kitzie et al. (2018) has produced the most recent bot-centric research covering mass shootings. It examined the retweet patterns and associated narratives of more than 400 social bot accounts—identified by submitting a random sample of total user accounts for classification via Botometer—which were active in the Parkland School Shooting Twitter conversation. Nied et al. (2017) conducted an exploratory analysis by hand-labelling suspected bots and examining alternative narratives—derived from the alternative narrative work of Starbird (2017)—in detected community clusters of OSN conversations of late 2015 discussing the Paris Attacks and the Umpqua Community College mass shooting. This particular paper significantly expands the current literature in terms of the scale of observed OSN conversations (i.e., 46.7 million tweets, 11.7 million Twitter user accounts), as well as directly comparing multiple event conversations.

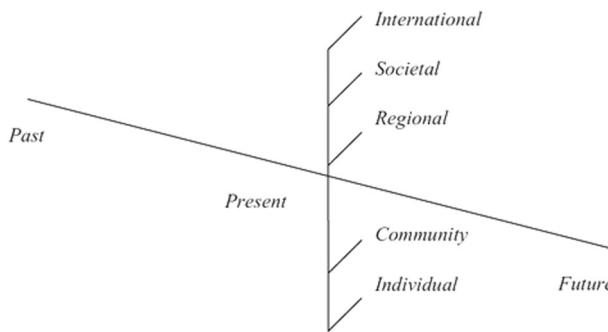


Fig. 1 Two-dimensional analytical framework of Chyi and McCombs (2004) for comparing frame changes across similar media events.

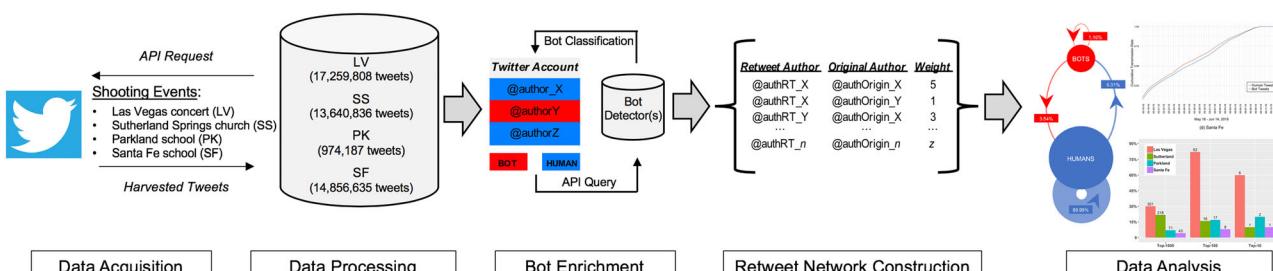


Fig. 2 Overview of social bot analysis framework illustrating methodological steps taken to analyse social bots within online social network conversations involving mass shooting events from October 2017 through May 2018.

Table 1 Summary of mass shooting events resulting in more than 10 deaths from October 1, 2017 through May 18, 2018.

Event	Date	Killed	Injured	Description
Las Vegas Concert Shooting (Las Vegas, Nevada)	Oct 1, 2017	59	418	1On October 1, 2017, gunman Stephen Paddock opened fire from his Las Vegas Mandalay Bay hotel room onto the Route 91 Country Musical Festival crowds killing 58 and himself, while injuring 418.
Sutherland Springs Church Shooting (Sutherland Springs, Texas)	Nov. 5, 2017	27	20	2On November 5, 2017, gunman Devin Kelley fired into the congregation of the Sutherland Springs First Baptist Church, killing 26 and himself, while injuring 20.
Parkland School Shooting (Parkland, Florida)	Feb. 14, 2018	17	17	3On February 14, 2018, gunman Nikolas Cruz opened fire within the Marjory Stoneman Douglas High School in Parkland killing 17, while injuring 17.
Santa Fe School Shooting (Santa Fe, Texas)	May 18, 2018	10	13	4On May 18, 2018, gunman Dimitrios Pagourtzis opened fire within the Santa Fe High School in Santa Fe killing 10, while injuring 13.

1. <https://www.nytimes.com/interactive/2018/10/01/us/las-vegas-shooting-victims.html>
 2. <https://www.cnn.com/2017/11/05/us/texas-church-shooting/index.html>
 3. <https://www.cbsnews.com/feature/parkland-florida-school-shooting/>
 4. <https://www.cnn.com/us/live-news/santa-fe-texas-shooting>

network construction' outlines the steps taken to create a network graph object of the retweet network for each OSN mass shooting conversation. Finally, 'Data analysis' introduces the methods employed to comparatively analyse the evidence of suspected bots across this study's OSN mass shooting conversations of interest.

Data acquisition and processing. Four mass shooting events that took place within an eight-month period from October 2017 through May 2018 serve as the mass shooting use-cases we analyse for this study. While additional shooting events meeting the generally accepted mass shooting threshold of at least four or more deaths in a single event (Towers et al., 2015; Dahmen et al., 2018; Silva and Capellan, 2019) occurred during this period, we deliberately focus on events that resulted in 10 or more deaths, since total victim counts serve as the most salient predictor of increased media interest and coverage (Schildkraut et al., 2018). We list these events in chronological order along with additional pertinent details in Table 1.

To derive the associated OSN conversations, we harvested streaming tweets via the available Twitter public application programming interface (API) for approximately a one-month period (28 days) following the date of each mass shooting event. In an effort to maintain a consistent collection paradigm for each event, we used Twitter API request parameters which relied on the same filter keywords: *shooting*, *shot*, *shots*, *gunman*, *gunfire*, *shooter* and *activeshooter*. The resulting corpus harvest for all four mass shooting events returned approximately 46.7 million tweets produced by approximately 11.7 million unique Twitter users. Table 2 provides volume metrics for each mass shooting event. We should note that the overall Parkland collection effort, although it employed the same search parameters and method, returned a substantially smaller total tweet volume due to the fact that the collection took place in a different collection environment with different collection parameters. Consequently, this specific dataset was then further filtered using the keywords indicated above, resulting in a smaller volume of data. Nevertheless, this data can still be considered as representing the Twitter narrative about the Parkland event for the purpose of this study (i.e., the relative ratios between tweets and retweets).

Bot enrichment. The open-source DeBot bot detection platform (Chavoshi et al., 2016), which enables retrospective analysis of historically detected bots through an archival repository, served as the primary source for labelling likely social bot accounts within our harvested tweet corpus. This is because the historical dates of our tweets were beyond the temporal classification constraint of the Botometer open-source bot detection platform (Davis et al., 2016). DeBot has proven to classify bots at extremely high precision rates in relation to other social bot detection efforts, including Twitter (Chavoshi et al., 2016, 2017). While DeBot classifies bots with great precision (i.e., greater true positive classifications), we acknowledge, in agreement with Morstatter et al. (2016), that such high precision comes at a cost to recall performance (i.e., greater false negative classifications).

The identification and labelling of social bot accounts follows a three-step process. First, the unique Twitter account name and numeric identification for each account present in the mass shooting event corpus are submitted for classification to the DeBot API. Next, DeBot returns a binary (True or False) bot classification for each user account. Finally, the DeBot classification results are merged with the existing corpus data by creating a true or false bot attribute for each account. In total, DeBot classified fewer than 1% of all corpus tweet account users, or contributors, as likely social bots responsible for producing ~1.63 million tweets and ~1.40 million retweets, or 3.49% and 3.92% of the tweets and retweets in

Table 2 Overall tweet corpus volumes and suspected social bot contributions for each associated OSN mass shooting event conversation.

Corpus	Collection dates	Tweets	Retweets	Contributors
Las Vegas	Oct 1-Oct 28, 2017	17,259,808	13,257,950	5,851,616
<i>Bot source (% of total)</i>		719,499 (4.17%)	622,610 (4.70%)	39,956 (0.68%)
Sutherland Springs	Nov 5-Dec 3, 2017	13,640,836	10,094,746	4,996,779
<i>Bot source (% of total)</i>		491,210 (3.60%)	418,771 (4.15%)	34,497 (0.69%)
Parkland	Feb 14-Mar 13, 2018	974,187	802,212	425,941
<i>Bot source (% of total)</i>		52,206 (5.36%)	45,507 (5.67%)	8,441 (1.98%)
Santa Fe	May 18-Jun 14, 2018	14,856,635	11,688,116	5,262,635
<i>Bot source (% of total)</i>		367,193 (2.47%)	316,368 (2.71%)	28,072 (0.53%)

the corpus, respectively. Table 2 provides applicable social bot volume details for each of the OSN conversations.

Retweet network construction. The deliberate act of retweeting has been viewed as an artefact demonstrating a particular Twitter user's propensity to share information or attempt to engage in direct conversation with other users (Boyd et al., 2010). Retweets accounted for 76.7% of the total mass shooting corpus, with approximately 35.8 million tweets identified as retweets. This overall high density of retweets permeates across each individual OSN event conversation, with retweet densities of 76.8%, 74.0%, 82.3% and 78.7% for the Las Vegas, Sutherland Springs, Parkland, and Santa Fe shooting conversations, respectively. A retweet between two Twitter accounts (i.e., nodes) is an observable conversational activity that can be viewed as a directed connection, or edge, within a network construct. For example, we assign a directed edge weight value of "1" for an initial retweet between two users and increment previously established edges by "1" for each subsequent directional retweet between the same two user nodes.

By iterating through each retweet in the corpus, we ultimately created a social network graph of each OSN mass shooting conversation. The transformation of conversations into network graph objects enables the application of an array of social network analysis (SNA) methods, which we detail in the subsequent 'Data analysis methods' section. Overall, the OSN retweet conversations produced directed networks with the following node-edge characteristics: 4,926,906 nodes/11,864,672 edges (Las Vegas), 4,105,206 nodes/8,987,800 edges (Sutherland Springs), 382,797 nodes/751,255 edges (Parkland) and 5,264,937 nodes/13,133,371 edges (Santa Fe).

Data analysis methods. The following subsections provide a comprehensive introduction to the specific methods employed to comparatively analyse the evidence of suspected social bots across the OSN mass shooting conversations of interest. Each subsection describes the fundamental data requirement for each analysis method, a detailed characterisation of the analysis method, and any pertinent theoretical underpinnings. The combined effort of these 'Data analysis methods' subsections provides necessary interpretative context to the presented findings in the subsequent 'Results and discussion' section of the study.

Conversation participation rate analysis. To determine any potential contribution patterns of social bot accounts in comparison to regular human accounts, we examined the cumulative tweet and retweet rates over the course of each observed online mass shooting conversation. We accomplished this by bifurcating each mass shooting event corpus into separate human and bot activity, then temporally indexing the contribution activity for each of these subsets. This allowed for a quantifiable and visual comparative analysis between human and bot temporal

conversation contributions to better understand if bot and human accounts differed in the rate of participation throughout the conversation collection period of analysis. An observable contribution rate difference between the two entities could serve as a potential temporal proxy of relative interest. We present the results of this comparative analysis and discuss the effects of bots as potential narrative drivers in the 'Conversation contribution inversion' subsection within the 'Results and discussion' section.

Analysis of subsequent mention of previous mass shooting events. Many traditional media studies have focused on comparatively analysing media attention between highly salient mass shooting events (e.g., Muschert and Carr, 2006; Schildkraut and Muschert, 2014; Schildkraut et al., 2018). To extend a comparative framework perspective to multiple OSN mass shooting conversations, we sought to determine the sustained attention paid to previous mass shooting events in subsequent mass shooting events. We accomplished this by observing the explicit mention rates of keywords associated with past events in subsequent events from both the human and suspected bot perspective. As previously explained, we maintained a consistent collection paradigm by using the same filter keywords to harvest the overall conversations for each mass shooting event but required event-specific keywords to determine specific mentions within other event conversations. We accomplished this by selecting the most common descriptive words within the corpus for each mass shooting event and using those emergent keywords to tally subsequent mentions in other events. The resulting mention keywords followed an "event name/shooter name" paradigm, as the most common distinguishable words for each previous event included a derivation of the specific event (Las Vegas {"vegas"}; Sutherland Springs {"sutherland"}; Parkland {"parkland"}) and reference to the identified shooter of each previous event (Las Vegas {"paddock"}; Sutherland Springs {"kelley"}; Parkland {"cruz", "nikolas"}). In the Parkland case, we had to differentiate between the shooter, Nikolas Cruz, and the Texan politician, Ted Cruz, so we added the first name "nikolas" for inclusion with "cruz" as a bigram for mining mention counts within the corpus. Figure 3 provides a visual framework describing the process we used for determining mention counts across events. We present the normalised mention rate results and discuss the perceived implications of previous mass shooting event mentions by humans and bot accounts in the 'Previous event mention rates' subsection of the 'Results and discussion' section.

Intra-group and cross-group interaction analysis. Having created retweet networks from the retweets of each OSN conversation, we sought to determine if the interactions *between* (i.e., social bots retweeting humans or humans retweeting social bots) and *among* (i.e., social bots retweeting social bots or humans retweeting humans) the different account types (i.e., humans or bots)

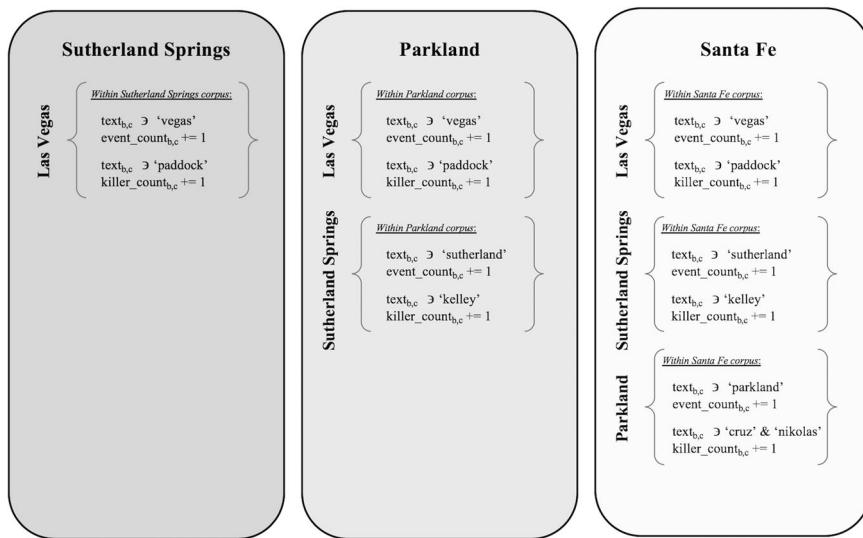


Fig. 3 Mention count discovery of previous mass shooting events within subsequent online mass shooting event conversations. For example, within the Santa Fe conversation corpus, a previous event mention count by both humans and bots is determined by references to event location (i.e., “vegas”, “sutherland”, or “parkland”) and event shooter (i.e., “paddock”, “kelley”, or “cruz and nikolas”).

Table 3 Retweet volumes of intra-group and cross-group conversation activity across all online mass shooting events.

Corpus	Intra-group retweet volume		Cross-group retweet volume	
	Bot-to-Bot	Human-to-Human	Human-to-Bot	Bot-to-Human
Las Vegas	153,468	11,931,552	704,064	469,152
Sutherland Springs	67,219	9,307,052	369,181	351,556
Parkland	4,552	711,027	45,692	40,956
Santa Fe	23,953	11,027,195	344,700	292,421

produced observable patterns. To do so, we had to subset each mass shooting event network into separate retweet network edgelists representing each potential cross-group and intra-group network edge relationship (i.e., bot-retweets-bot, bot-retweets-human, human-retweets-bot, human-retweets-human). Table 3 presents the consolidated volumes associated with all derived edgelist relationships for each online mass shooting event. These edgelist volumes serve as the basis for the presented intra-group and cross-group conversation rate results detailed in ‘Intra-group and cross-group conversation patterns’ subsection of the ‘Results and discussion’ section.

Relative importance of conversation contributors through centrality analysis. We take further advantage of the graph constructed for each derived retweet network through the application of SNA centrality measures. Centrality measures serve as a proxy to determine the relative importance of a node based on a given node’s structural network position vis-a-vis other nodes (Wasserman and Faust, 1994). Riquelme, González-Cantergiani (2016) presented a comprehensive survey examining the wide variety of available centrality measurements that can be used to measure the relative influence of contributing users in online Twitter conversations. In an effort to determine the relative importance of social bot actors in relation to human actors, we chose to calculate the following four network centrality

measurements based on their recognisability and efficient scalability to large-scale networks for all nodes within each of the mass shooting retweet networks: in-degree centrality, out-degree centrality, eigenvector centrality, and PageRank centrality.

In-degree and out-degree centrality are directional variants of degree centrality, which simply measures the total number of direct edges that a node shares with other nodes in a given network. In the context of a retweet network, in-degree centrality provides a cumulative inward activity tally of all inbound edges to a particular node, or the number of times a message from a particular Twitter account is retweeted by nodes in the network. The reverse is true of out-degree, as it measures all outward activity of a particular node, or the number of times a particular Twitter account initiates a retweet of other messages produced by nodes in the network. Measures of degree centrality can be viewed as a proxy for network popularity given the quantifiable number of direct connections, or conversation engagements. Eigenvector centrality, a more complex derivation of degree centrality, queries the individual degree centrality of all nodes in a network and returns a weighted sum based on a particular node’s set of direct and indirect edges (Bonacich, 2007). Given the completeness of the eigenvector calculation across an entire network, we can view eigenvector centrality as a measure of global network influence in a retweet network. Lastly, the PageRank centrality measurement, derived from eigenvector centrality, places a weighted premium on the degree value of nodes that initiate edges with other nodes of the most relative importance (Brin and Page, 1998). From a retweet network perspective, user accounts with higher PageRank valuation receive more retweets from the more popular user accounts in the retweet network. The subsection ‘Relative importance of social bots in online mass shooting conversations’ within the ‘Results and discussion’ section presents our formal centrality analysis results and discusses the overall density of social bots within the highest ranking centrality measures.

Results and discussion

In the following, we present and discuss the results of the applied methods described in the previous ‘Data and methods’ section. Through the acquisition of a Twitter data corpus from multiple online mass shooting conversations and the identification of

social bots within this corpus, we were able to apply the previously described analysis methods to comparatively analyse social bot conversation participation in relation to human user behaviour and present our findings in the subsequent 'Conversation contribution inversion' and the 'Previous event mention rates' subsections. Furthermore, through the application of SNA techniques to the mass shooting event retweet networks, we present our findings of directional conversation interactions in the 'Intra-group and cross-group conversational patterns' subsection and the centrality analysis rankings of bots in relation to humans in the 'Relative importance of social bots in online mass shooting conversations' subsection.

Conversation contribution inversion. Figure 4 provides a consolidated visualisation of the cumulative contribution profiles for human and suspected bot accounts over the course of each mass shooting conversation. We observed that the pace of cumulative tweet contributions from bots exceeds that of humans through the entire conversation timeframe for the Sutherland Springs (Fig. 4b) and Santa Fe (Fig. 4d) shootings. However, with the Las Vegas (Fig. 4a) and Parkland (Fig. 4c) shootings, an inversion occurs as bots begin to outpace humans after five and seven days (annotated as grey shaded areas), respectively, and continue on as the leading contributor in terms of a normalised contribution rate.

Persistent human contribution latency in relation to bots for the entirety of the Sutherland Springs and Santa Fe conversations suggests that human accounts lacked general interest compared to bots for these particular events. In contrast, the Las Vegas and Parkland events draw immediate human interest, but this initial interest subsides as sustained bot contribution rates bypass humans after less than a week. This discovery will be discussed in more detail in the further works section. However, in discovering the clear contribution rate differences between bots and humans in the Sutherland Springs and Santa Fe events and, more interestingly, the contribution inversion in the Las Vegas and Parkland events, we can conclude that social bots are an explicit sub-population of actors in online mass shooting event conversations. Further, in the same light that Merry (2016) identified special interest groups as narrative framing agents in social media, social bots should be considered as potential actors capable of framing online narratives.

Previous event mention rates. By executing the mention count discovery process introduced and illustrated (Fig. 3) in the 'Data analysis methods' section, we were able to ascertain the rates at which bots and humans mentioned previous events in the subsequent mass shooting events of this study's corpus. Within this paradigm, we observe historical event mention relationships as follows: Las Vegas mentions within the Sutherland Springs, Parkland and Santa Fe conversations; Sutherland Springs mentions within the Parkland and Santa Fe conversations; Parkland mentions within the Santa Fe conversation. Table 4 presents the consolidated total mention volumes and mention rates of the Las Vegas, Sutherland Springs and Parkland mass shootings within the applicable subsequent mass shooting conversations. We normalised the mention rates according to the unique bot and human populations within each conversation. The results show a clear partiality by both bots and humans towards mentioning associated event names (i.e., "vegas", "sutherland", "parkland") in lieu of the identified shooter (i.e., "paddock", "kelley", "cruz") when discussing previous mass shooting events. Furthermore, we observe no mentions of Devin Kelley, the Sutherland Springs gunman, in any other mass shooting event conversations. Finally, while Sutherland Springs struggles to garner any attention by the

time of the Santa Fe conversation, we see a drastic increase in both human and bot mention rates of Las Vegas. While there is little research available to properly classify these observable patterns from an OSN-specific view, we can look to traditional media studies to potentially contextualise these findings. For example, Levin and Wiest (2018) discovered that media consumers paid significantly more attention to shooting events when the narrative focused on courageous bystanders as opposed to a victim or killer, while Silva and Capellan (2019) presented an extensive overview of observable media attention patterns in mass shooting media research.

Intra-group and cross-group conversation patterns. Figure 5 presents an overall visualisation of normalised intra-group (i.e., bots retweeting bots or humans retweeting humans) and cross-group (i.e., bots retweeting humans or humans retweeting bots) conversation patterns for each of the study's OSN conversations. We normalised the intra-group and cross-group retweet volumes by the total retweet volume in each conversation. For example, the self-loops for humans and bots depicted in Fig. 5a translate to 89.99% of all the retweets in the Las Vegas corpus resulting from human-to-human interaction, while bot-to-bot retweets only account for 1.16% of retweets. In addition, the directed cross-group activity between bots and humans in the Las Vegas conversation shows that humans retweeting bots comprises 5.31% of retweets, while bots retweeting humans comprises 3.54% of retweets. In general, we see human-to-human interaction as the dominant relationship across all of the mass shooting conversations. Moreover, human users retweet bots at a higher rate than bots retweet humans in each of the conversations, which demonstrates that humans are more responsible for spreading bot-generated content than bots themselves in each of the mass shooting conversations. This is an interesting finding, as previous social bot analysis has found bots to be more, on average, hyper-social than humans: they attempt to engage humans at persistently higher rates in retweet networks associated with election, conflict, and political Twitter conversations as opposed to average human accounts (Stella et al., 2018; Schuchard et al., 2019).

In extending the analysis beyond aggregate interaction volume rates, Table 5 presents the range of contribution volumes by user type (i.e., bot or human) across all conversations. We see consistent median user contribution volumes for bots that are three-fold to four-fold higher than human accounts. Corresponding range values from both the tweet and retweet volume are also annotated to determine accounts with extremely high contribution volumes. Interestingly, top human accounts contribute far more total tweets than the top bot accounts from an aggregate tweet perspective for each conversation. However, from a retweet subset perspective, the highest social bot contributor in the Las Vegas and Santa Fe retweet conversations far exceeded the highest human contributor. These resulting super contributors warrant additional analysis and should serve as primary points of evaluation for the future work content analysis which is proposed in the conclusion of this paper.

Relative importance of social bots in online mass shooting conversations. The results of the centrality measurement ranking analysis, introduced in the 'Data and methods' section, showed that many social bots, while accounting for just 0.67% of all contributors in this study's corpus, displayed structural network importance by achieving high centrality ranking positions, especially in the eigenvector and out-degree centrality rankings. Figure 6 presents the consolidated centrality results, depicting the density of social bot accounts falling within the top- N , where $N = 1000/100/10$, eigenvector, in-degree, out-degree and PageRank centrality rankings for

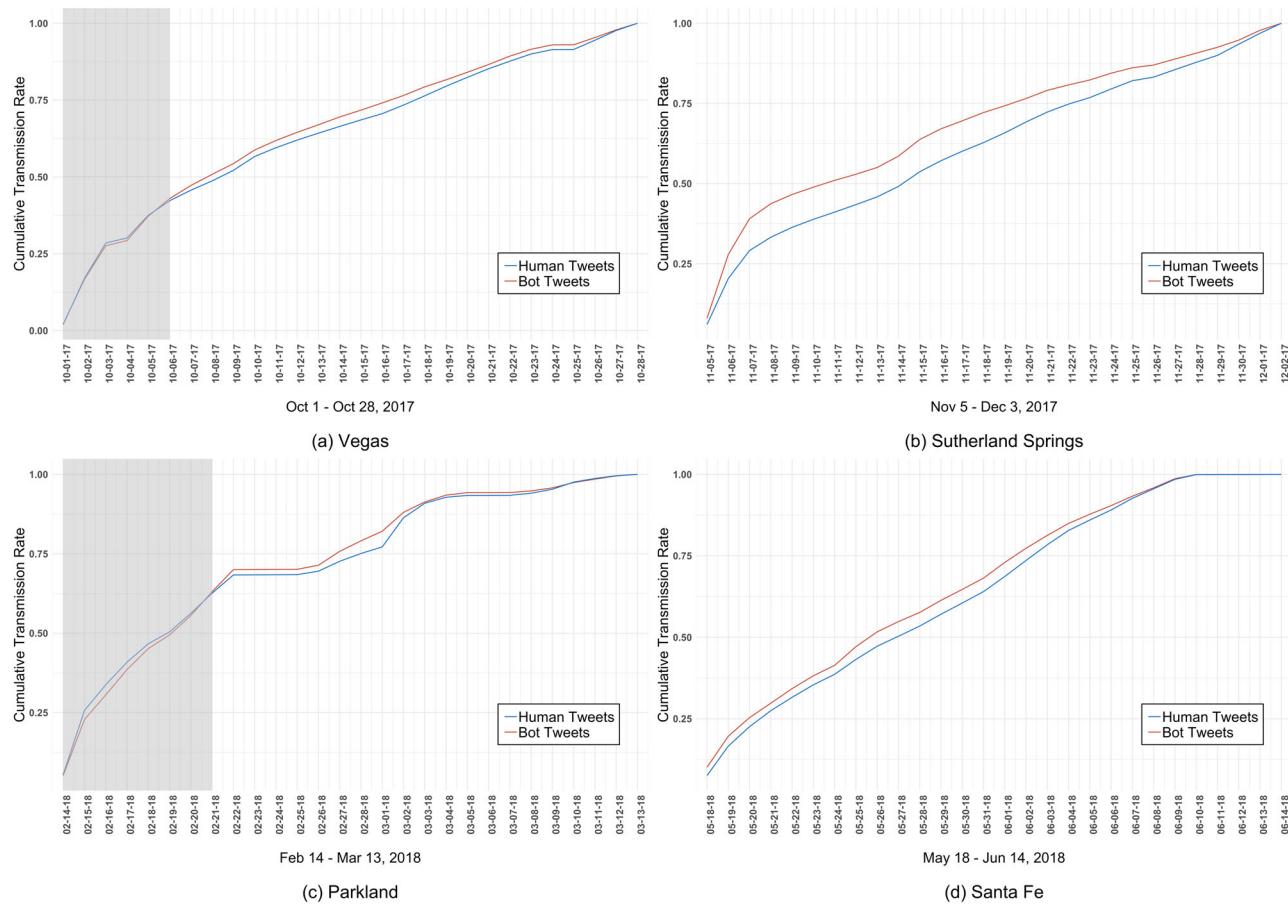


Fig. 4 Cumulative tweet conversation contributions of both human (blue) and bot (red) accounts for the one-month online conversations of the following mass shooting events: **(a)** Las Vegas concert shooting (October 1-28, 2017), **(b)** Sutherland Springs church shooting (November 5-December 3, 2017), **(c)** Parkland school shooting (February 14-March 13, 2018), **(d)** Santa Fe school shooting (May 18-June 14, 2018). Grey shaded areas depict human-led contribution rate periods.

Table 4 Bot and human mention rates of previous mass shooting events in subsequent mass shooting conversations.

SUTHERLAND CONVERSATION				PARKLAND CONVERSATION				SANTA FE CONVERSATION							
Bots	34,497			Bots	8,441			Bots	28,072						
Humans	4,962,282			Humans	417,500			Humans	5,234,563						
		Mentions	Rate			Mentions	Rate			Mentions	Rate				
'vegas'	Bot	12,431	0.36035	'vegas'	Bot	191	0.02263	'vegas'	Bot	3,860	0.13750				
	Human	165,927	0.03344		Human	2,362	0.00566		Human	71,016	0.01357				
'paddock'	Bot	1,031	0.02989	'paddock'	Bot	0	0	'paddock'	Bot	367	0.01307				
	Human	9,858	0.00199		Human	0	0		Human	2,723	0.00052				
Total Las Vegas Mentions		Mentions	Rate	Total Las Vegas Mentions		Mentions	Rate	Total Las Vegas Mentions		Mentions	Rate				
Bot		13,462	0.39024	Bot		191	0.02263	Bot		4,227	0.15058				
Human		175,785	0.03542	Human		2,362	0.00566	Human		73,739	0.01409				
						'sutherland'	Bot	47	0.00557	'sutherland'	Bot	226	0.00805		
							Human	674	0.00161		Human	3,399	0.00065		
						'kelley'	Bot	0	0	'kelley'	Bot	0	0		
							Human	0	0		Human	0	0		
						Total Sutherland Mentions	Bot	47	0.00557	Total Sutherland Mentions		Mentions	Rate		
							Human	674	0.00161	Bot		226	0.00805		
									Human		3,399	0.00065			
										'parkland'	Bot	8,334	0.29688		
											Human	141,176	0.02697		
										'cruz'	Bot	426	0.01518		
											Human	5,664	0.00108		
						Total Parkland Mentions	Bot			Mentions	Rate				
							Human			8,760	0.31205				
												146,840	0.02805		

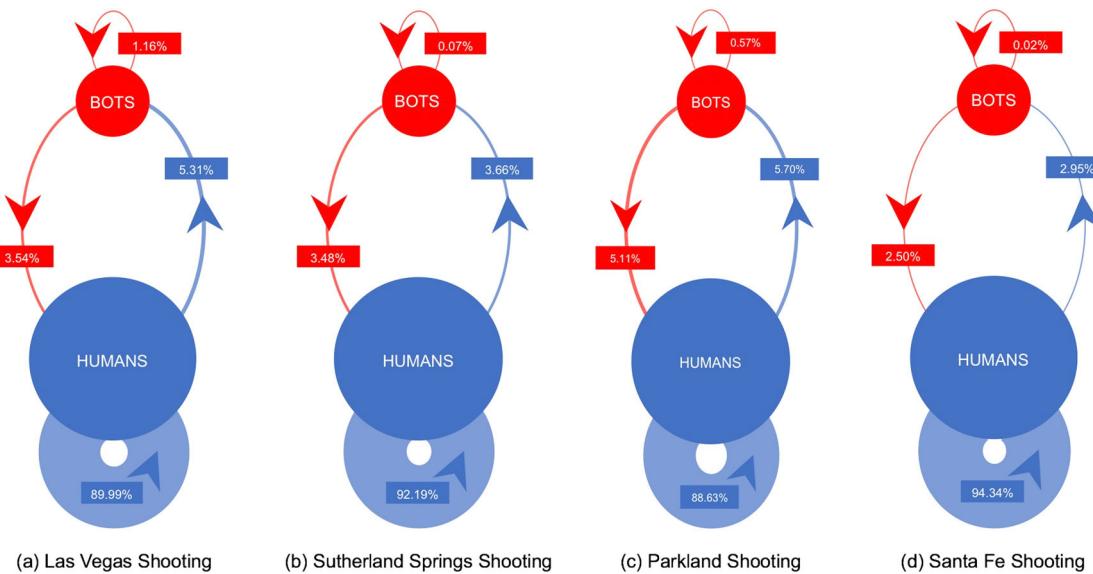


Fig. 5 Intra-group and cross-group retweet interaction rates among and between human (blue) and suspected social bot (red) user accounts for a one-month period following the (a) Las Vegas, (b) Sutherland Springs, (c) Parkland and (d) Santa Fe shooting events.

Table 5 User type contribution volumes within each OSN conversation corpus of interest.

Corpus	User type	Median user Tweet volume	Min/max user tweet volume (# of users)	Median user retweet volume	Min/max user retweet volume (# of users)
Las Vegas shooting	Humans	1.0	Min: 1 (3,424,165 users) Max: 6,713 (1 user)	1.0	Min: 1 (2,796,826 users) Max: 1,737 (1 user)
	Bots	4.0	Min: 1 (10,021 users) Max: 3,530 (1 user)	4.0	Min: 1 (9,654 users) Max: 3,338 (1 user)
Sutherland Springs shooting	Humans	1.0	Min: 1 (3,041,764 users) Max: 6,683 (1 user)	1.0	Min: 1 (2,420,661 users) Max: 1,536 (1 user)
	Bots	4.0	Min: 1 (9,070 users) Max: 2,292 (1 user)	4.0	Min: 1 (8,772 users) Max: 641 (1 user)
Parkland shooting	Humans	1.0	Min: 1 (284,272 users) Max: 1,214 (1 user)	1.0	Min: 1 (248,731 users) Max: 945 (1 user)
	Bots	3.0	Min: 1 (2,802 users) Max: 662 (1 user)	3.0	Min: 1 (2,750 users) Max: 413 (1 user)
Santa Fe shooting	Humans	1.0	Min: 1 (3,113,882 users) Max: 5,173 (1 user)	1.0	Min: 1 (2,547,297 users) Max: 2,673 (1 user)
	Bots	4.0	Min: 1 (7,429 users) Max: 3,586 (1 user)	4.0	Min: 1 (7,172 users) Max: 3,559 (1 user)

each online mass shooting conversation. The out-degree ranking persistence shows the hyper-social nature of these particular social bots across all conversations. More interestingly, social bots display exceedingly high eigenvector centrality valuations in the Las Vegas conversation, accounting for 82% and 60% of the top-100 and top-10 rankings, respectively, while also earning numerous high rankings across the other conversations, but not nearly as dominant. Given that eigenvector centrality serves as a potential proxy for total network influence, social bots could be construed as the most structurally influential nodes in the Las Vegas mass shooting conversation. In contrast to the high in-degree and eigenvector ranking results, the sparse in-degree and PageRank centrality ranking results suggest that social bots are not initial attractors of attention, but rather achieve their structural relevance by their persistent actions to reach out to other accounts.

Conclusion

In this study, we examined the presence, contribution patterns, and relative importance of suspected social bots within four

different online mass shooting conversations. This work contributes to the greater understanding of OSNs as a primary news source and the ability of them, along with new content providers (i.e., social bots), to frame a sustained, participative conversation surrounding the potentially contentious narrative focused on mass shootings. By following a mixed methodology process focused on the normalisation and fusion of associated Twitter conversation data with social bot detection results, we presented a repeatable and agnostic process that can be extended to evaluate additional online mass shooting use cases of interest. While analysing the cumulative contribution patterns of both humans and bots, this study found that social bot accounts outpaced human contributions throughout the entirety of the Sutherland Springs and Santa Fe shooting conversations. With the Las Vegas and Parkland conversations, however, a reversal took place. Although human accounts initially outpaced bot accounts during the first week, bots outpaced humans for the remainder of the conversations. Both bots and humans displayed a strong tendency to mention previous shooting events by referencing event locations as opposed to shooters themselves. The construction of

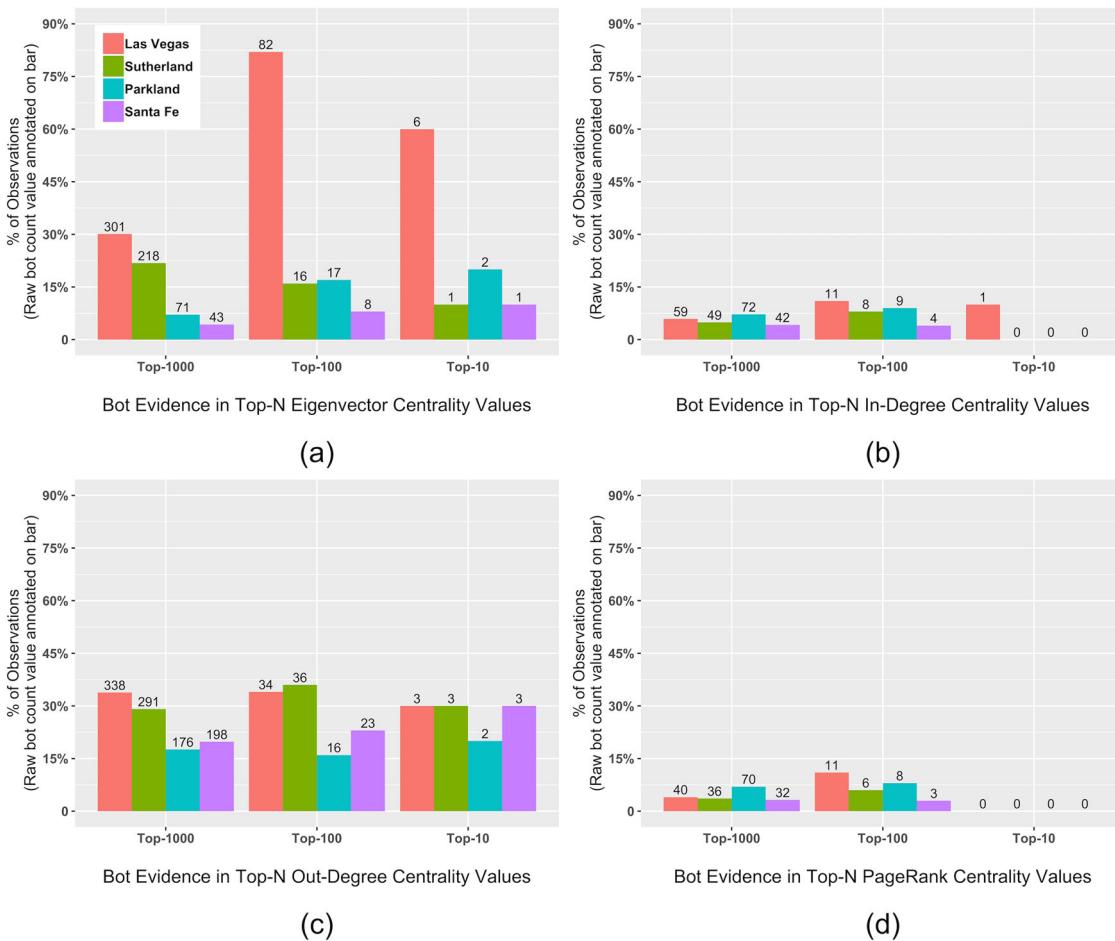


Fig. 6 Social bot accounts in the top-N, where $N = 1000/100/10$, (a) eigenvector, (b) in-degree, (c) out-degree and (d) PageRank centrality measurement rankings within OSN mass shooting retweet networks discussing the Las Vegas (red), Sutherland Springs (green), Parkland (blue) and Santa Fe (purple) shooti.

retweet networks allowed us to observe the intra-group and cross-group engagements showing humans engaging bot accounts at a higher rate than bots engaging humans in all four conversation retweet networks. Moreover, the retweet network graph construct enabled the application of SNA centrality measurements to investigate the relative importance of social bots, which showed large populations of bots ranking prominently in overall eigenvector and out-degree centrality across all conversations. In the Las Vegas mass shooting conversation, social bots dominated the eigenvector centrality ranking results, accounting for 82% and 60% of the top-100 and top-10 accounts, respectively.

This study is not immune from limitations. First, previous traditional media research efforts have shown significant bias in mass shooting media coverage based on event factors such as the race/ethnicity of both the shooter and the victims, as well as the number of associated casualties (Duxbury et al., 2018; Schildkraut et al., 2018). Guggenheim et al. (2015) described the reciprocal relationship between traditional and social media discussing mass shooting events and we can assume the transference of this coverage bias. In other works, Tufekci (2014) points to well-known inherent biases associated with data emanating from OSNs, to include sampling and representativeness. In addition, Ruths and Pfeffer (2014) echoed a similar critique of OSN data, while also stressing the inability of OSN providers to prevent or limit the distortion of social bot actors.

Further work could potentially overcome these limitations by incorporating additional OSN platform sources (e.g., Facebook,

Instagram) and expanding the analysis to additional mass shooting events that unfortunately continue to take place. This research focused on the identification and relative importance of social bots within the observed conversations, but additional analysis could seek to gain insights into the design motivations of specific bots based on the content they produce beyond this study's identification of previous shooter name and shooter event mention rates across subsequent events. Through the application of deeper natural language processing techniques, a deliberate content analysis methodology could seek to determine the extent to which specific social bots amplify a particular narrative. Further, the introduction of community detection analysis could identify the existence of particular bots aligning with verifiable political beliefs of other community members within the conversations. Finally, the conversation inversion discovery provides an opportunity to further delve into why identified bot accounts display more sustained interest in shooting events over time than human accounts. One possible explanation could be that the interest of human accounts evolves with the reframing of the mass shooting discussion, while bot accounts maintain their focus on the initial presentation framing of the mass shooting event.

The unprecedented scale of this study serves as a compelling first step forward in providing the social bot analysis research necessary to identify and distinguish automated social bot contributions from human dialogue in OSNs discussing mass shootings. Given the evidence of contagion in the aftermath of mass shootings (Towers et al., 2015), it is essential to detect and

prevent the potential amplification or glorification of such events by social bots. While social bot analysis is still in a nascent state and bot detection methodologies continue to evolve to account for the growing sophistication of bot developers (Subrahmanian et al., 2016; Cresci et al., 2017), this work provides a repeatable framework that is extendable to other OSN conversations and additional bot detection platforms. Further, it provides requisite feedback to bot detection algorithm developers on associated detection performance against an array of different OSN conversations. Finally, this research highlights the essential need to understand the role that social bots play as a potential emergent mode to address the opinions and attitudes that people have about firearms and firearm policies in online conversations about mass shootings. However, while this study focused on just a subset of recent mass shootings, it is imperative that concerted research expand rapidly into this area in an effort to help stem the tide of these events that continue to wreak havoc upon the lives of innocent individuals.

Data availability

The datasets generated during and/or analysed during the current study are available in the Harvard Dataverse repository: <https://dataverse.harvard.edu/privateurl.xhtml?token=5f8bd221-15f9-405b-85c6-04811878cadd>.

Received: 14 August 2019; Accepted: 28 October 2019;

Published online: 10 December 2019

Notes

1 Botometer is accessible at <https://botometer.iuni.iu.edu>

2 DeBot is accessible at <https://www.cs.umn.edu/~chavoshi/debot/>

References

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. Preprint at arXiv:14078134

Bessi A, Ferrara E (2016) Social bots distort the 2016 U.S. Presidential election online discussion. *First Monday* 21

Bolsover G, Howard P (2017) Computational propaganda and political big data: moving toward a more critical research agenda. *Big Data* 5:273–276. <https://doi.org/10.1089/big.2017.29024.cpr>

Bonacich P (2007) Some unique properties of eigenvector centrality. *Soc Networks* 29:555–564. <https://doi.org/10.1016/j.socnet.2007.04.002>

Boshmaf Y, Muslukhov I, Beznosov K, Ripeanu M (2013) Design and analysis of a social botnet. *Comput Networks* 57:556–578. <https://doi.org/10.1016/j.comnet.2012.06.006>

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, DC, pp 1–10

Brin S, Page L (1998) The anatomy of a large-scale hypertextual Web search engine. *Comput Networks ISDN Syst* 30:107–117. [https://doi.org/10.1016/S0169-7552\(98\)00110-X](https://doi.org/10.1016/S0169-7552(98)00110-X)

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

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: *Proceedings of the 26th International Conference on World Wide Web Companion*. pp 1601–1606

Chong D, Druckman JN (2007) Framing theory. *Annu Rev Political Sci* 10:103–126. <https://doi.org/10.1146/annurev.polisci.10.072805.103054>

Chyi HI, McCombs M (2004) Media salience and the process of framing: coverage of the Columbine school shootings. *Journal Mass Commun Quart* 81:22–35. <https://doi.org/10.1177/107769900408100103>

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, Switzerland, pp 963–972

Dahmen NS, Abdenour J, McIntyre K, Noga-Styron KE (2018) Covering mass shootings: Journalists' perceptions of coverage and factors influencing attitudes. *Journal Practice* 12:456–476

Davis CA, Varol O, Ferrara E, et al. (2016) BotOrNot: a system to evaluate social bots. In: *Proceedings of the 25th International Conference Companion on World Wide Web*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, pp 273–274

Duxbury SW, Frizzell LC, Lindsay SL (2018) Mental illness, the media, and the moral politics of mass violence: The role of race in mass shootings coverage. *J Res Crime Delinquency* 55:766–797

Edwards F, Howard PN, Joyce M (2013) Digital activism and non-violent conflict. *SSRN* 2595115

Entman RM (1993) Framing: toward clarification of a fractured paradigm. *J Commun* 43:51–58. <https://doi.org/10.1111/j.1460-2466.1993.tb01304.x>

Ferrara E, Varol O, Davis C, et al. (2016) The rise of social bots. *Commun ACM* 59:96–104. <https://doi.org/10.1145/2818717>

Guggenheim L, Jang SM, Bae SY, Neuman WR (2015) The dynamics of issue frame competition in traditional and social media. *Ann Am Acad Political Soc Sci* 659:207–224. <https://doi.org/10.1177/0002716215570549>

Howard PN, Kollanyi B (2016) Bots, #StrongerIn, and #Brexit: Computational Propaganda during the UK-EU Referendum. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.2798311>

Kitzke VL, Mohammadi E, Karami A (2018) "Life never matters in the DEMOCRATS MIND": Examining strategies of retweeted social bots during a mass shooting event. *Proc Assoc Inform Sci Technol* 55:254–263. <https://doi.org/10.1002/pra2.2018.14505501028>

Lazer DMJ, Baum MA, Benkler Y, et al. (2018) The science of fake news. *Science* 359:1094–1096. <https://doi.org/10.1126/science.aaq2998>

Levin J, Wiest JB (2018) Covering mass murder: an experimental examination of the effect of news focus—killer, victim, or hero—on reader interest. *Am Behav Scientist* 62:181–194

Mahabir R, Croitoru A, Crooks A, et al. (2018) News coverage, digital activism, and geographical saliency: A case study of refugee camps and volunteered geographical information. *PLoS ONE* 13:e0206825. <https://doi.org/10.1371/journal.pone.0206825>

Merry MK (2016) Constructing policy narratives in 140 characters or less: the case of gun policy organizations. *Policy Studies Journal* 44:373–395. <https://doi.org/10.1111/psj.12142>

Mitchell A (2018) Americans still prefer watching to reading the news—and mostly still through television. *Pew Research Center*

Moffat BS (2019) Medical Response to Mass Shootings. In: Lynn M, Lieberman H, Lynn L, et al. (eds) *Disasters and mass casualty incidents: The nuts and bolts of preparedness and response to protracted and sudden onset emergencies*. Springer International Publishing, Cham, pp 71–74

Mønsted B, Sapiezyński P, Ferrara E, Lehmann S (2017) Evidence of complex contagion of information in social media: an experiment using Twitter bots. *PLoS ONE* 12:e0184148. <https://doi.org/10.1371/journal.pone.0184148>

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

Muschert GW, Carr D (2006) Media salience and frame changing across events: coverage of nine school shootings, 1997–2001. *Journal Mass Commun Quart* 83:747–766

Newman BJ, Hartman TK (2017) Mass shootings and public support for gun control. *Br J Political Sci* 1–27

Nied AC, Stewart L, Spiro E, Starbird K (2017) Alternative narratives of crisis events: communities and social botnets engaged on social media. In: *Companion of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. ACM, New York, pp 263–266

Riquelme F, González-Cantergiani P (2016) Measuring user influence on Twitter: a survey. *Inform Proc Manag* 52:949–975. <https://doi.org/10.1016/j.ipm.2016.04.003>

Ruths D, Pfeffer J (2014) Social media for large studies of behavior. *Science* 346:1063–1064. <https://doi.org/10.1126/science.346.6213.1063>

Schildkraut J, Elsass HJ, Meredith K (2018) Mass shootings and the media: why all events are not created equal. *J Crime Justice* 41:223–243

Schildkraut J, Muschert GW (2014) Media salience and the framing of mass murder in schools: a comparison of the Columbine and Sandy Hook massacres. *Homicide Stud* 18:23–43

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, 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. <https://doi.org/10.1038/s41467-018-06930-7>

Silva JR, Capellan JA (2019) The media's coverage of mass public shootings in America: fifty years of newsworthiness. *Int J Comparative Appl Criminal Justice* 43:77–97

Starbird K (2017) Examining the alternative media ecosystem through the production of alternative narratives of mass shooting events on Twitter. In: Eleventh International AAAI Conference on Web and Social Media

Stella M, Ferrara E, Domenico MD (2018) Bots increase exposure to negative and inflammatory content in online social systems. *PNAS* 115:12435–12440. <https://doi.org/10.1073/pnas.1803470115>

Subrahmanian VS, Azaria A, Durst S, et al. (2016) The DARPA Twitter bot challenge. *Computer* 49:38–46. <https://doi.org/10.1109/MC.2016.183>

Towers S, Gomez-Lievano A, Khan M, et al. (2015) Contagion in mass killings and school shootings. *PLoS ONE* 10:e0117259. <https://doi.org/10.1371/journal.pone.0117259>

Tufekci Z (2014) Big Questions for social media big data: representativeness, validity and other methodological pitfalls. *Ann Arbor, Michigan*, pp 505–514

Varol O, Ferrara E, Davis CA, et al. (2017) Online human-bot interactions: detection, estimation, and characterization. In: Eleventh international AAAI conference on web and social media

Vosoughi S, Roy D, Aral S (2018) The spread of true and false news online. *Science* 359:1146–1151. <https://doi.org/10.1126/science.aap9559>

Wasserman S, Faust K (1994) Social network. Analysis: methods and applications, 1 edn. Cambridge University Press, Cambridge, New York

Acknowledgements

Special appreciation to the DeBot team at the University of New Mexico led by Nikan Chavoshi for access to the DeBot platform and discussing its applicability for this use case. A.S.'s contributions based upon work partially supported by the U.S. Department of Homeland Security under a COE Grant Award Number 2017-ST-061-CINA01. The

views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the U.S. Department of Homeland Security.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to R.S.

Reprints and permission information is available at <http://www.nature.com/reprints>

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



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2019