Terrorism and Political Violence

ISSN: (Print) (Online) Journal homepage: www.tandfonline.com/journals/ftpv20

Automating Terror: The Role and Impact of Telegram Bots in the Islamic State's Online Ecosystem

Abdullah Alrhmoun, Charlie Winter & János Kertész

To cite this article: Abdullah Alrhmoun, Charlie Winter & János Kertész (2024) Automating Terror: The Role and Impact of Telegram Bots in the Islamic State's Online Ecosystem, Terrorism and Political Violence, 36:4, 409-424, DOI: 10.1080/09546553.2023.2169141

To link to this article: https://doi.org/10.1080/09546553.2023.2169141

TERRORISM AND POLITICAL VIOLENCE
2024, VOL. 36, NO. 4, 409–424
https://doi.org/10.1080/09546553.2023.2169141

Routledge
Taylor & Francis Group

Automating Terror: The Role and Impact of Telegram Bots in the Islamic State's Online Ecosystem

Abdullah Alrhmoun [a,b], Charlie Winter[c], and János Kertész [a,d]

[a]Department of Network and Data Science, Central European University, Vienna, Austria; [b]Observatory on Social Media, Indiana University, Bloomington, Indiana, USA; [c]ExTrac AI, London, UK; [d]Complexity Science Hub, Vienna, Austria

ABSTRACT

In this article, we use network science to explore the topology of the Islamic State's "terrorist bot" network on the online social media platform Telegram, empirically identifying its connections to the Islamic State supporter-run groups and channels that operate across the platform, with which these bots form bipartite structures. As part of this, we examine the diverse activities of the bots to determine the extent to which they operate in synchrony with one another as well as explore their impacts. We show that these bots are mainly clustered around two communities of Islamic State supporters, or "munasirun," with one community focusing on facilitating discussion and exchange, and the other one augmenting content distribution efforts. Operating as such, this network of bots is used to lubricate and augment the Islamic State's influence activities, including facilitating content amplification and community cultivation efforts, and connecting people with the movement based on common behaviors, shared interests, and/or ideological proximity while minimizing risk for the broader organization.

KEYWORDS

Terrorist bots; online extremism; bot communities

Introduction

To understand and ultimately mitigate the threat from terrorism today, which is increasingly reliant on the massive, multivariate usage of internet-based technologies, researchers, policymakers, and practitioners require new approaches that rely on complexity science.[1] Over the past decade, network theory has contributed tremendously to our understanding of how and why violent extremist communities form and thrive.[2] Among other things, network science has helped to demonstrate that the topological structure of these illicit networks has common features with that of other complex systems and social phenomena,[3] even though extremist communities are often treated as a social aberration at a policy level.[4] At the same time, features specific to terrorist networks have been revealed through network science, like the counter-intuitive role women play in making them robust,[5] and the relationship between structural network characteristics and the severity of the attacks carried out by the actor in question.[6]

In recent years in particular, several valuable efforts have been made to develop and/or apply quantitative tools to identify and analyze complex structures present within violent extremist networks in online spaces. Among these tools are community detection algorithms,[7] which can be used to find topologically related clusters reflecting similar interests and activities in groups of nodes in social ecosystems. Using these algorithms along with a range of other approaches, several scholars have attempted to unearth and understand community structures within terrorist networks, tracing and anticipating the dynamics that shape them and locating common characteristics among their members.[8]

CONTACT Abdullah Alrhmoun alrhmoun_abdullah@phd.ceu.edu Department of Network and Data Science, Central European University, Quellenstrasse 51, 3rd floor, Vienna 1100, Austria

This article has been corrected with minor changes. These changes do not impact the academic content of the article.

Supplemental data for this article can be accessed online at https://doi.org/10.1080/09546553.2023.2169141.

Today, one of the most important but under-researched aspects of terrorist networking online is the use of bots, something that we explore in this article using network science methods. Specifically, we investigate the role of bots as they appear in the context of Islamic State supporter communities on Telegram. An instant messaging app and social media platform, Telegram is favored by the Islamic State,[9] (as well as many other violent extremist movements) on account of its broad array of functionalities—which includes anything from content hosting and broadcasting to peer-to-peer chats—and branding focus on user privacy/user sovereignty.[10] Although most of Telegram's offerings are currently unencrypted, users can choose to implement encryption technologies when using its services. Besides its peer-to-peer messaging functionality, Telegram has channels and groups. Channels are one-way broadcasting lists, where only admins can send messages and users can subscribe (and unsubscribe). They can be of two types: public and private. Anyone can join a public channel, but users need an invitation link, sometimes shared publicly, to join private channels. In groups, which can also be public or private, anyone can send a message and interact with other members of the group, whether they are an admin or not.

Bots on social media, including but not limited to Telegram, are best understood as automated accounts that execute specific tasks such as publishing, sharing, and resharing content.[11] Anyone can build and deploy a bot: All that is required is a sustainable and open Application Programming Interface (API) access. In recent years, bot networks have been used increasingly across both mainstream and lesser known platforms (including Telegram) to promote products and services,[12] spread misinformation,[13] and low-credibility content,[14] probe algorithmic and political bias,[15] manipulate elections and public opinions,[16] and mitigate the challenges posed by violent extremism.[17] As this study demonstrates, in the context of the Islamic State and its supporters' activities on Telegram, bots generally perform one of three key functions: publishing content, moderating discussions, and acting as gatekeepers. In this capacity, they play a central, lubricating role in amplifying the movement's ideology and cultivating its community of sympathizers, automating administrative tasks like blocking users that violate group policies, and permitting new members to join.

Below, drawing on 1,215,850 data points that were collected from Telegram between February 1 and September 30, 2021 via an ingest program, we study the activities of the Islamic State's "terrorist bots" within the community dynamics amidst which they operate. We map out their interaction network and, in addition to presenting a schema of their activities and impacts, we sequentially apply community detection algorithms to the data with a view to determining the extent to which they operate in a structured or unstructured manner. These methodologically distinct approaches parse the network's underlying structure by dividing its nodes into communities. All of the applied methods show that the structure is by no means randomly distributed but, rather, made up of clusters (or modules) of closely interconnected nodes. Our analysis of the network's modular structure and clustered activities implies the existence of a hybrid system of functional groupings that have been proactively, and collectively, developed to augment the Islamic State's presence in Telegram channels and groups, as well as a spontaneous process of unorganized supporter-generated community formation. Based on these findings—which speak to the flexibility, ease, and effectiveness with which bots can be deployed to further the interests of bad actors—we contend that the allowances that Telegram makes for bot development are a central factor driving the Islamic State's years-long preference for it over other platforms that are demonstrably more secure. This is in spite of the (valid and widely implemented) assertion by Telegram in its FAQs that "we do block terrorist (e.g., ISIS-related) bots and channels."[18]

The article proceeds as follows. First, we give a brief overview of the literature on violent extremist activism online, focusing in particular on Islamist militancy and, more specifically, jihadism. Next, we set out the data collection and analytical methodology, explaining how we collected the data on the bot network and how we processed and interrogated it. After that, we describe the overarching characteristics of the network, touching on what functions these bots performed, how frequently they were active and for how long, and in what language they operated. We then present the findings of the community detection analysis itself; in it, we explore the clustered and utility-driven topology of the network. We conclude by weighing up the implications of this study and suggesting further avenues for research.

Literature review

Scholars from a wide array of disciplines have long tracked how online terrorist activities have developed in synergy with advances in technology and shifts in both the physical and information security environment. Collectively, this work has demonstrated that, much of the time, developmental trends are intuitive and borne of an iterative process of bottom-up innovation.[19] However, occasionally—and often when the shifts they result in are most impactful—there is also evidence of top-down influence playing a role.[20]

For decades, Islamist violent extremists have set out to adopt and exploit new technologies to facilitate their operations (both military/terroristic and recruitment-focused). As Torres-Soriano has pointed out, the late 1990s and early 2000s saw prominent organizations like the Global Islamic Media Front (GIMF) and Al Qaeda transitioning from static websites to closed forums.[21] After more than a decade of use, these were, in the first half of the 2010s, then supplanted by "conventional" social media platforms like Twitter and Facebook.[22] Following effective and systematic targeted disruption from these mainstream tools, most jihadist outreach online is now confined to the partially encrypted broadcasting and chat platform Telegram, which, while it has certainly become increasingly inhospitable to militancy-related activism in recent years, remains a more open and functional space than sites like Twitter.[23]

In recent years, numerous researchers have pointed out that, despite its more optimal security and functionality, Telegram does not have a full monopoly on jihadist outreach. Other, similarly orientated and secured apps like WhatsApp, Element and Hoop have also emerged as preferred platforms for sensitive communication between adherents of extremist Islamism.[24] Moreover, since 2018 in particular, static websites have become increasingly important spaces once more, especially in the context of propaganda archiving and distribution. And, as Winter, Sayed and Alrhmoun have observed, for more conventional state-based groups like Hamas, Hizbullah and the Afghan Taliban, Twitter has never been more important.[25] That being said, Telegram still remains a hegemonic presence for jihadist outreach activities online, activities that typically take one of two overlapping forms: propaganda production and distribution, and group identity formation. As Wagemakers observes, jihadist organizations have long invested a significant amount of their time and energy in propaganda work.[26] From the influence network of Abu Jandal al-Azdi, one of Al Qaeda's (AQ) most important Internet ideologues in the late 2010s, to al-Shabab's sprawling covert media apparatus, online spaces have long been replete with examples of jihadist organizational outreach.

Importantly, while it is to date arguably the most prominent example of jihadist propagandizing, the Islamic State is by no means the only militant Islamist group to have seen the value in resource-intensive strategic outreach. Hamas, Hizbullah, and the Afghan Taliban all preside over similarly sized (if not bigger) media networks, which, as Khatib notes, they each deploy—like the Islamic State—to shape the narrative landscape and attract followers, legitimize their actions, and intimidate their opponents.[27] Moreover, alongside their official, organizational outreach infrastructures—whether in the context of Sunni or Shi'i militancy—a vast array of supporter-run media outlets and agencies operate, peddling their respective ideological line and amplifying their messages—often, incidentally, with assistance from automated bots.[28]

In the more tangible, operationally direct sphere of identity formation (of which active recruitment can be a downstream consequence), jihadist outreach is characterized by a combination of hierarchical design and organic, volunteer-led activism. The role of social media platforms—Twitter in particular—in social absorption and formal enlistment has been pivotal over the course of the last decade, as indicated in numerous case study-led assessments focusing on foreign fighter networks in Syria (including several network analyses).[29] These studies have shown that on-/offline recruitment networks, while largely organic, are often carefully groomed spaces populated by official operatives and unofficial advocates, with the latter serving as connectors or beacons that directly elicit engagement from curious onlookers, drawing them in before furnishing them with the information they need to physically sign up—i.e., who to talk to, where to fly to, how to evade being apprehended, and so on.[30]

While these online encounters are important, however, a measure of face-to-face interaction is usually also required to facilitate the process of joining any militant movement, something that accounts for the continuing prevalence of what Conway describes as real world, social network-based recruitment patterns.[31]

Recognizing the new centrality of Telegram to extremist networking online, in the last four years in particular, scholars have begun to turn away from Twitter to focus their attention on how these same community behaviors and activities pan out on other platforms like Telegram, prominent among them Clifford and Powell and, separately, Amarasingam, Maher, and Winter.[32] However, to date, there have been no quantitative or network-based analyses of extremist networks on this platform, presumably down to the fact that there are significantly more technical obstacles when it comes to collecting Telegram data versus Twitter data (issues that are set out below in the Methodology section). This leaves us with a partial understanding of how and why Telegram continues to be a preferred arena for supporters of groups like the Islamic State today—a gap in the knowledge that this article seeks to go some of the way towards remedying.

Moreover, while there have been a number of exploratory and speculative studies looking at the role of bots in terrorist networks, these have so far refrained from addressing the specific challenges they present on Telegram. Among the first to substantively trace the practical impact of bots in the context of terrorism was Berger, whose early analysis of the role of the Islamic State's "Dawn of the Glad Tidings" app on Twitter was something of a trailblazer.[33] Berger showed that, through this app, the Islamic State was able to post 40,000 synchronized tweets in a single day when Mosul fell to its forces in 2014. While not technically driven by bot activity, this effort was an early precursor to the group's broader experimentation with artificially augmented social media activities.[34] Other analyses from the likes of Bondy in 2017 and Sultan in 2019 have been more forwards-looking, speculating as to how bots might be deployed to support malign activities at both a state and non-state level in the course of the next decade. While useful, these and studies like them are largely hypothetical, drawing on anecdotal observations or small datasets. On that basis, in this article, we hope to both add to and elevate the existing knowledge base regarding terrorist usage(s) of bots.[35]

Methodology

This section gives an overview of both the data collection process and the analytical methodologies that were used to interrogate said data.

Collecting the data

In order to better understand the role and impacts of bots deployed in support of the Islamic State on Telegram, we collected an original dataset from a manually selected community of 3,940 public groups and channels that were considered to be either controlled by or supportive of the Islamic State. The data collection period was February to September 2021. We selected these groups and channels for analysis because they were explicitly and proactively aligned with the Islamic State. A group or channel was deemed to be "explicitly aligned" if, on a sustained basis, its users' focus was on either news or ideological matters relating to the Islamic State. Other overt indicators of pro-Islamic State orientation that were taken into account include: re-posting or sharing of official media, the creation of unofficial pro-Islamic State media, or overt declarations of support by group administrators for the movement's mission, goals, activities, and operations.[36] Once selected for inclusion, each group or channel was tagged according to the topic it prioritized. In total, we applied sixteen tags to the dataset, meaning that the groups and channels generally revolved around one of sixteen issues or spheres of activity. These are: general commentary; news; Afghanistan; anti-Shi'a; al-Hol/Roj; Kashmir/India; information security; links sharing; media campaigns; medical advice; nashids; content distribution; networking; theology; tactics, techniques, and procedures (TTPs); and women's affairs. The largest category by far was "general commentary," which describes groups and channels (but mainly the former) in which

anything from conflict news to theology is discussed and both unofficial and official propaganda materials are shared. It is important to state that this is not a complete sample of all Islamic State channels and groups on Telegram, however, it does serve as an extensive representation of the pro-Islamic State community on the platform.

Cleaning the data

Once we had selected this sample, we accessed and archived all the publicly available messages shared by these groups and channels—programmatically stripping away all metadata besides the unique (and randomized) identification numbers associated with accounts that published on them. Next, we requested user objects from Telegram's API using the methods "users.getFullUser" and "users.getUsers." This provided us with the usernames of all the active accounts in the network that had opted to make their user details listed on Telegram's public directory when signing up to the platform. In total, we requested the usernames of 16,682 accounts through Telegram's API, recognizing that the vast majority of them would not be publicly listed—and, indeed, only 258 were. Determining which of these 258 publicly listed accounts were bots and which were not was simple, because Telegram mandates that all bots have usernames that end with the word "bot" (such as tetris_bot or "TetrisBot").[37] This meant we could filter the user data to include only information related to or messages posted by bots as well as identifying all posts that contained a URL referring to bot.

In total, we identified 106 accounts in this network that were bots. Having ascertained this, we discarded all other user data and collected all names and usernames directly associated with these bots, including those of the groups and channels in which they had been most active. We also collected all dates and times of posts and relevant data about content type (i.e., was it a JPEG, MP3, MP4, or PDF). While some of these bots are Telegram natives, the majority of them were compiled by self-appointed Islamic State khuruq al-buwatt ("bot creators"), who advertise their bot development services as part of their support for the broader extremist community. (Interestingly, these self-described khuruq typically label bots as hudhud, which is a reference to the Islamic scripture; the hudhud was a bird which served under the Prophet Solomon and supported him by collecting information about his enemies).[38]

Constructing the network

In order to determine how, and to what end, these bots were deployed by Islamic State supporters during the data collection period, we built a network map capturing their actions and interactions in the ecosystem. The nodes of this bipartite network, which is visualized in Figure 1, are bots (blue) and channels and groups (green and yellow respectively). Although this network is bipartite, the platform architecture is multilevel, due to the ability to link channels to groups, where each new post from a channel can automatically be forwarded to a connected discussion group by a so-called "discuss link."[39] The edges are directed links representing messages posted by bots in said groups/channels or mentions of these bots in said groups/channels. Although the edges represent two types of linkage, we opted to treat them as one for the sake of simplicity (given the unwieldy nature of network maps that are both bipartite and multilayer). We placed a link between a given bot and group or channel if the bot was found to have posted one or more messages in the group or channel in question. When this happens, it indicates that the bot is a group member or channel admin. Bots usually post far more than one message, but, for the purposes of this aspect of the study, we opted to ignore the weight of the link and count only the existence of the connection, not its strength.

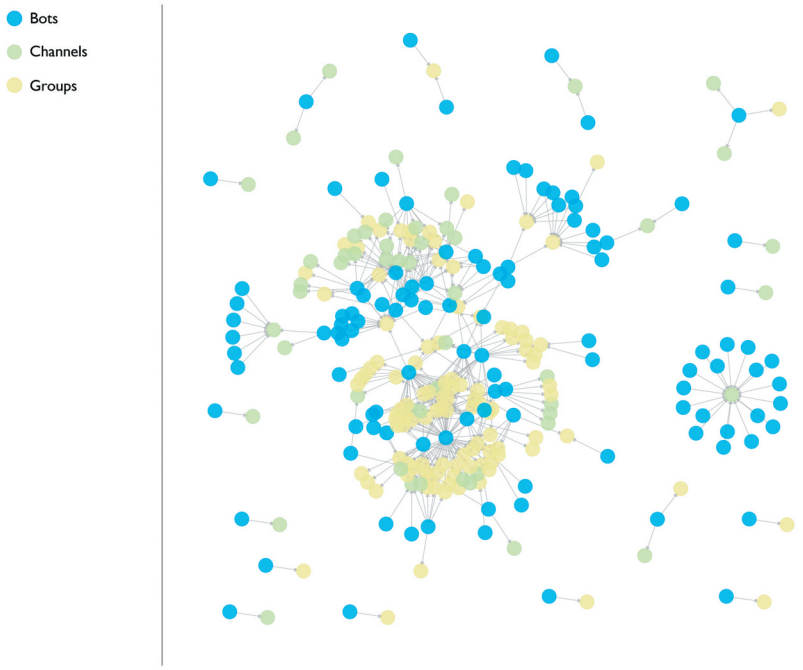

Figure 1. The visualization of the bots-groups-channels network. The blue nodes are bots, the green nodes are channels, and the yellow nodes are groups. The network consists of 241 nodes and 346 links. 106 of the nodes are bots, ninety-eight are groups, and thirty-seven are channels.

Identifying communities

In network science, communities refer to parts of a connected network within which nodes are linked to each other more than they are to the rest of the network. Such communities are generally understood to have a functional role and their efficient identification is a major scientific challenge.[40] Communities can be disjunct—in which case finding them means to identify a partition of the network—or they can be overlapping, leading to a so-called cover of the network.[41] Considering the multitude of tasks carried out by bots on Telegram, it is a natural question to ask whether, in this network of bots, channels and groups, we can identify whether the implementation of these tasks is organized, or whether its topology is decoupled from the respective functions of the bots. To study this problem, we used four different community detection methods to parse the network. We present and compare the outcomes of this analysis in the Appendix. While the methods yield somewhat different communities, common features emerge indicating robust topological structures. To better illustrate these features, we use the partition resulting from the Girvan-Newman (GN) algorithm[42] (see Figure 4), which represents them in the clearest way and leads to the most meaningful community structure. (The GN algorithm is known to be a powerful method, though only when restricted to relatively small networks, which makes it ideal for the present context.)

Findings

The 106 bots we identified via Telegram's API were either active in or mentioned on ninety-eight distinct groups and thirty-seven distinct channels. Their level of activity was significant: collectively, they posted some 39,211 messages. Most of these were published in groups (33,487 in groups and 5,724 in channels).

The vast majority of the content the bots published was text (36,715 posts of the 39,211 we collected). After that was image files (1,748), PDFs (315), video clips (185), audio files (ninety), and

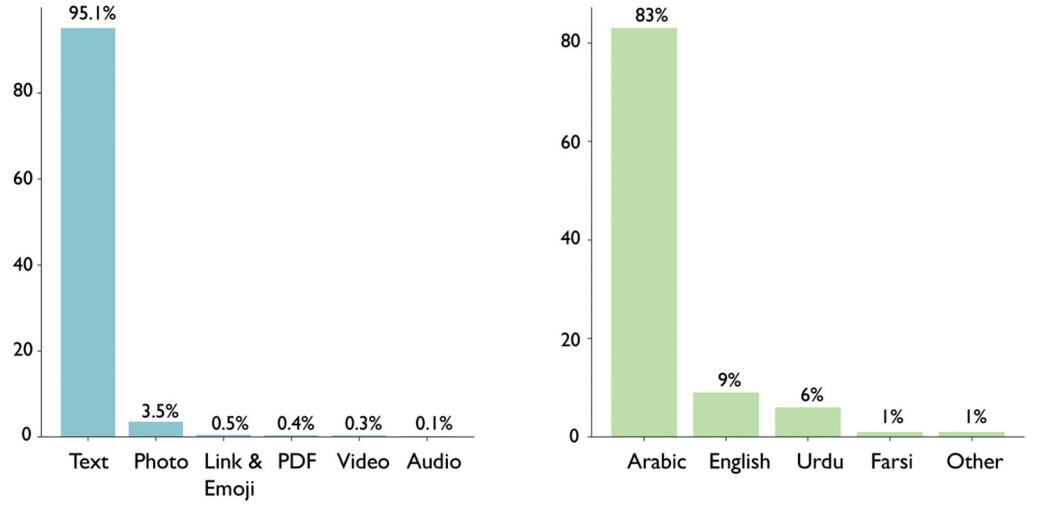

Figure 2. Left: The percentage of content types published by the bots. Most of the content is text; the rest comprises of image files, PDF files, video clips, audio files, and RAR files and emojis. Right: The language distribution of the content. Most of the text content is Arabic, then English, Urdu, Farsi.

emojis and links (158). From a linguistic perspective, the bots mainly posted in Arabic (29,605 messages in total). Besides that, there were 4,194 posts in English, 2,123 in Urdu, 922 in Farsi, 147 in Soko, and forty-five in Bahasa-Indonesia, and a smattering in other languages. Figure 2 shows the percentage of the type of content and the top languages.

The bots' respective lifespan was highly variable. By the end of the data collection period, none of the non-Telegram native bots were active anymore, with some having been censored by Telegram and others having been shut down by their administrators. The longest bot lifespan was 213 days, and the shortest was just one day. On average, they were active for around eighteen days. Naively, one would expect a linear dependence between the bots' lifespan and their cumulative activity. However, Figure 3 shows that there are a lot of fluctuations in the activity, resulting in a correlation coefficient R = 0.65. Moreover, the median and mode of the lifespans of bots are both about one day, which suggests that the lifespan of individual bots is typically short but, due to some bots having longer periods of activity, the system as a whole is surprisingly robust. Given what we know about Telegram's disruption policy when it comes to the Islamic State, this suggests that longer-life bots engaged in acts that were less explicitly or overtly supportive of the movement (i.e., activities other than media distribution).

Notably, there were several instances in the data of the removal-recreation cycle that characterizes how pro-Islamic State communities respond to Telegram's moderation efforts. For example, the *@UrduNashir_22bot* was removed and recreated half a dozen times during the data collection period, each time coming back with a different but still recognizable username (e.g., *@UrduNashir_24bot, @UrduNashir_27bot* and *@Urdu_nashir28bot*). This means that, generally, the removal of individual bots did not explicitly cause sustained disruption to the network, which was able to function in spite of Telegram's censors.

On average, the bots posted 176 messages each day (140 text, thirteen images, twelve PDFs, four video clips, and two audio files). For reasons that are not immediately clear, they were most active on August 2, 2021, when they collectively posted 1,431 messages. By contrast, on a not insignificant number of days, they were almost completely inactive, posting just one message per day during such periods.

From the perspective of the range of their activities, some bots were specialized and others were multifunctional. Most bots were associated with "general commentary" channels and groups; of the eighty-six that operated in that context, there were twenty media amplification bots, twenty-four links sharing bots, nineteen media activism bots, and seventeen news-posting bots.

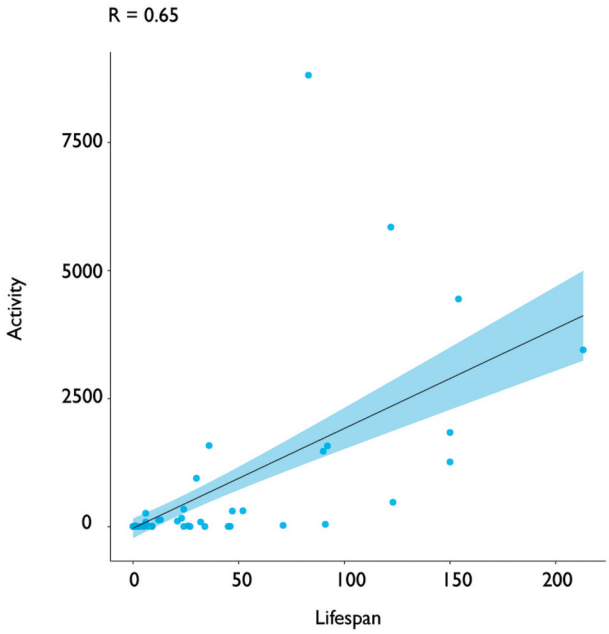

Figure 3. Pearson correlation coefficient R shows a positive relationship between the lifespan of the bots and their activity.

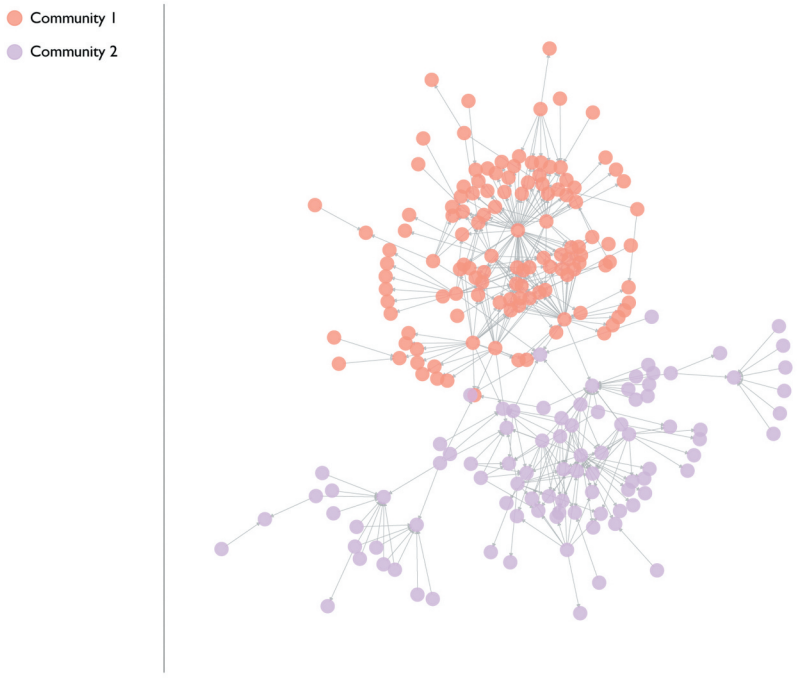

Figure 4. Visualization of the communities in the largest connected component of the bots-channels-groups network, as produced by the GN algorithm. It was partitioned into two parts with communities of similar sizes ninety-three (Orange) and 125 (pink).

Analyzing the communities

The bots-groups-channels network in question, as visualized above in Figure 1, consists of several components: small ones, consisting of two or three nodes; one medium size one, the shape of a star; and one large connected component containing the majority of the nodes. Below, we restrict our analysis to this largest component, which dominates the network. It is not clear *a priori* what governing principles could be behind this community structure. One possibility could be that partition would lead to its splitting into clusters of bots and clusters of groups/channels. This is because, at base, it is a bipartite and nebulous network (i.e., there are no direct connections among the groups and channels or the bots). Another possibility could be that the bots would get positioned randomly in the matrix of channels and groups, irrespective of their functions. On applying community detection algorithms to the data, however, we found that the communities we detected are in fact diverse and different but highly structured, containing an array of groups, channels and bots. Indeed, once parsed by the GN algorithm, we see two well separated communities that make up the structure of the network (see Figure 4).

One hundred and twenty-five nodes belong to Community 1, marked in orange, and ninety-three belong to Community 2, marked in pink (see Figure 4). Interestingly, while they are similarly sized, the composition of these two communities is rather different: forty-four nodes in Community 2 are groups/channels and forty-nine are bots; on the other hand, twenty-seven nodes in Community 1 are bots, and ninety-eight are groups/channels.

Activity analysis

Across the two communities, we identified two types of bots, with three core functions emerge: content distribution, basic group administration such as blocking spammers and deleting messages, and gatekeeping (i.e., allowing users to join and sharing links). Admin bots mostly perform administrative functions, including discourse moderation, link sharing, and gate keeping. Content bots engage directly with groups and channels, sharing content and exchanging information with group members. This latter grouping is clustered into many different clusters, with different bots connected with different channels or groups that publish different content types.

We manually tagged each bot as either being content-focused or admin-focused. Some bots can be both, but we tagged them by the majority of their activities. Community 1 has twenty admin bots and seven content bots and Community 2 has forty-two content bots and seven admin bots (see Figure 5). Notably, the number of channels/groups in Community 1 is twice as big (n = 98) as the number of channels/groups in Community 2 (n = 44). By reviewing the content that the bots published in each community, we were able to prove the validity of the overall partitioning process. That is to say, the two communities identified by the algorithm differed significantly when it comes to the amount and type of content they each posted. Community 1 published some 30,251 items, with Community 2 publishing just 6,284 items, even though it is only moderately smaller in scale. Notably, the content shared by Community 2 is significantly more diverse. Specifically, it shared 4,021 text-based posts, 1,562 images, 315 PDF files, 168 video clips, 158 links and emojis, and sixty audio files. In contrast, Community 1 overwhelmingly published text-based items; in total, it shared 30,073 text-based posts and just 135 images, twenty-nine audio files, and fourteen video clips. (See Figure 6). This suggests that Community 1 is focused more on facilitating/moderating chatter between group members than it is sharing content (official or otherwise), something that is broadly left to Community 2.

The two communities also differ linguistically. Figure 6 shows that the language of the text corpus in each community differs significantly. The bots in Community 1 overwhelmingly published Arabic messages (26,533), and significantly less in English (2,831). In contrast, the bots in Community 2 published messages in Urdu (1,781), in Arabic (1,328), in Farsi (823), and in English (486).

Each and every bot, group, and/or channel in these two communities serves a specific purpose for the broader ecosystem. In both communities, we found that their operations were trifurcated in a way that

Figure 5. Content bots versus admin bots. The graph shows that Community 1 has more admin bots than content bots and the admins are connected to more channels than in Community 2, which has significantly more content bots than admin bots.

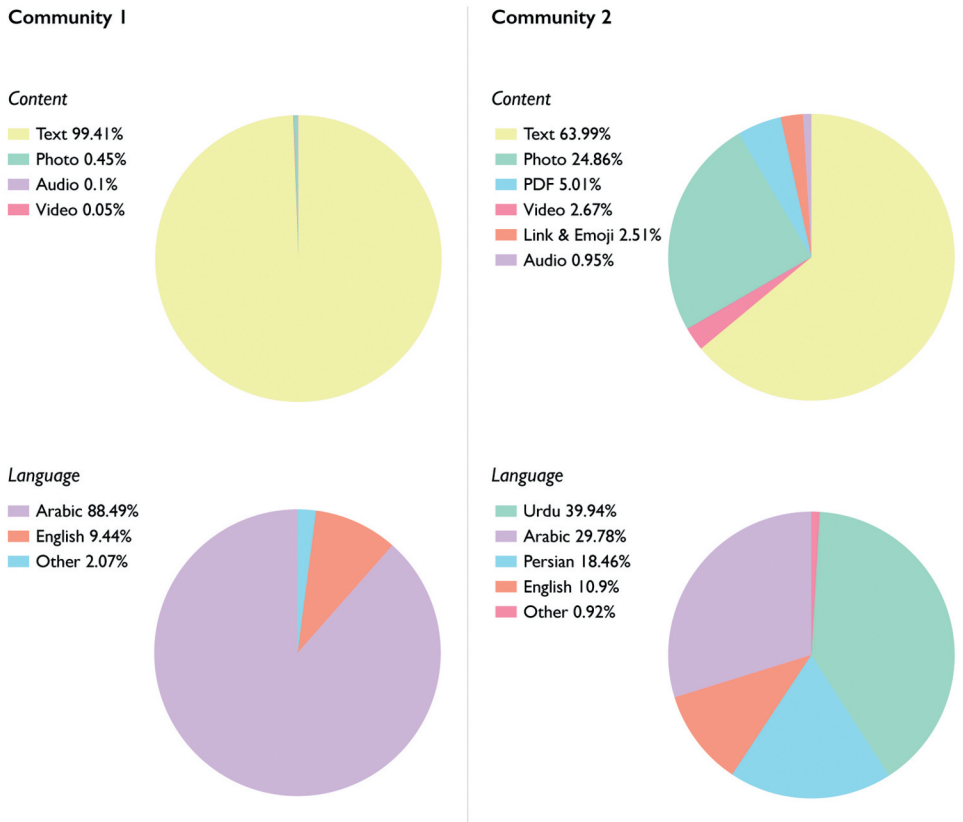

Figure 6. Type and language of content shared by Communities 1 and 2. It shows that Arabic text dominates Community 1 while Community 2 has more diverse content types and mixed languages.

spoke to both parallel missions and targeted audiences: general commentary on the Islamic State, often grounded in distribution of and discussion around content produced by its official media apparatus; media materials that have been translated and/or engineered to respond to/support the interests of local pro-Islamic State communities in specific locations or contexts (like ISKP outreach networks in Afghanistan and IDP camps in Syria and Iraq); and efforts geared towards facilitating the activities of the Islamic State and its supporters online (like information security advice and links to specific technical services).

Importantly, while they are structurally different, the two communities we identified do not exist in two entirely separate spaces. They engage with the same audience segments, but, due to their functionality and nature, have evolved to serve different but complementary purposes. That being said, their overarching audience and orientation is the same—aiding and abetting people who share and support Daesh's mission either online or offline, locally or internationally. The efficient and clearly resilient recreation cycle that characterises the Islamic State's presence on Telegram today, not to mention the persistent presence of bots within it, speaks to their widespread utility and impact.

Network structure analysis

Demonstrably, the network is not organized by single person or controlling authority but formed incrementally, at least partly in a self-organized manner, by loosely connected groups of Islamic State supporters. This corresponds to the broader decentralized character of the Islamic State's influence activities, a network quality that has been cultivated because it makes exceptionally difficult to localize and break down the ecosystem in a sustained and effective manner.

To identify the mesoscale structure of both communities and the position/role of the bots and channels/groups in this structure, we used Rombach's Algorithm to analyze the core-periphery (C-P) structure of both networks.[43] The algorithm is a modified version of the continuous Borgatti-Everett algorithm,[44] an improvement to detect multiple cores in the core-periphery structure in networks (See Appendix for the mathematical details of the algorithm). Simply put, the structure comprises a few core nodes and many periphery nodes, where the core is composed of densely interconnected nodes, and the periphery comprises sparsely interconnected nodes.

The communities exhibit C-P structure as a result of a partly spontaneous process (see Figure 7). The coreness of both communities is very close (the coreness of Community 1 and Community 2 is ≈ 0.30 and 0.29 respectively). Community 1′s core nodes are made up of many admin bots in the core. Indeed, the node with highest coreness is an administrator bot. This reflects our observation that Community 1 has the task of assuring information flow. Community 2, by contrast, has only one admin bot at the center of its core; the rest of its core nodes are content bots, which implies, as per our initial observation, that this community's main role is content distribution.

The different core-periphery structures and, as part of that, the proportion of admin bots over content bots, speaks to differences in how bots operate when they are focused on information flow over content distribution. The former involves keeping the network available (and "clean") to facilitate the latter. These bots are focused on making the content supply chain efficient; they do not produce or provide the goods, but they help to keep the distribution running (it is about control). The latter is all about production and provision, sharing content with as many members of the community as possible (it is about participation). Inasmuch as that is the case, and although this splitting of responsibilities is not clean-cut, the two communities share a common objective while having distinct roles.

In any case, the core of both communities is mainly comprised of bots, with their peripheries made up of a mixture of channels and groups. Based on that structure, one could argue that the core-periphery structures emerged because of constant activity from the core in maintaining content distribution efforts and chat moderation, which was being conducted in the periphery, assuming that the core–core relations are not as important as the core–periphery relations.

The network has two main attributes: (i) core-like nodes comprising admin bots that control and moderate the activities of groups and channels; and (ii) facilitating general commentaries and exchanges between Islamic State supporters, with less emphasis on content distribution. These

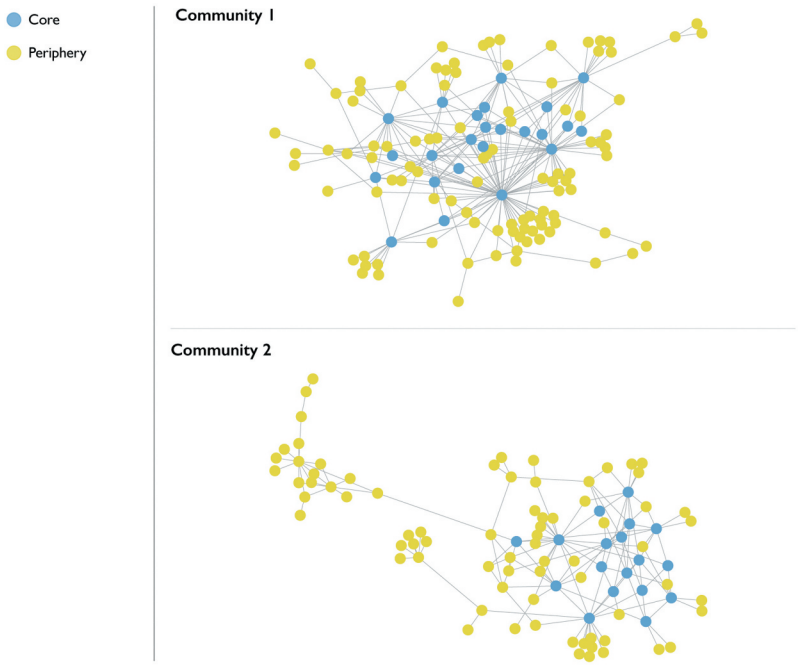

Figure 7. Network visualization that indicates the core and periphery nodes in both communities.

attributes together imply that its principal impact (and, consequently, utility) corresponds to its ability to maintain the activities of the broader ecosystem of supporters.

The contrasting structure of each network suggests that different bots serve different purposes and that, in doing so, their roles and applications have iteratively evolved alongside the communities they serve. In other words, the way these bots (or groups of bots) operate and/or are applied has helped to shape these networks in a manner that best serves their specific purposes. In community 1, one core can be seen dominating the network, which shows a level of discipline; the other community, on the other hand, has multiple cores, which speaks to a measure of flexibility and scope for equal engagement.

The evolution of these two communities, which has left one highly centralised and the other more multipolar, has a rough corollary in the Islamic State's own evolution as a political movement in recent years. Its shift from territorial protostate to covert insurgency in its core territories of Iraq and Syria, in tandem with sustained pressure from the online platforms it once favoured, forced the munasirin community to transform from a model of organic but acute centralisation to one characterized by multipolarity. That is to say, in its early days, both the Islamic State and its community of supporters prioritized discipline and control online. However, with its decline and ultimate territorial collapse, discipline and control were no longer feasible and were instead swapped for an approach that was more in line with the characterisation the second community we detected.

Importantly, these bots must be understood as being part of a bigger picture, not the picture itself; they alone do not constitute the ecosystem. Rather, they help to construct and curate its constituent networks (thus communities) and—by pushing the limits of scalability of human-made networks—they serve to augment and strengthen its operations. Inasmuch as this is the case, the bot network is a mirror of the larger human community within which it operates.

We can be confident, given the conditions within which it exists, that both the bot network and the human community have developed evolutionarily, responding opportunistically to the stresses, strains, and obstacles they have faced over the years. Due, however, to a lack of historical data, it is not possible to empirically trace the scale of that evolution—at least, not as part of this study. That being said, the communities in this network are not completely disconnected. On the contrary, they

are very much interlinked. However, each community's respective nodes are more connected with each other than they are with the nodes of the other community. On that basis, we believe these two groupings have evolved out of one original network entity.

Conclusion

On Telegram, anyone can create and manage a bot. This means that users can deploy bots to do most of the activities that human users can: post content, receive messages, manage groups, provide services, and even accept payments. Unlike human users, though, bots can be programmed to be active, and immediately responsive, around the clock. This makes them ideal for tasks that are otherwise labor-intensive or tedious. Indeed, they are in many ways better-placed to manage groups and share content because they can do so at large scale and with immediacy.

Bots can be used for good and bad; they can be truly harmful when used by violent extremists. Indeed, as we have shown above, for the Islamic State—similar to many other violent extremist groups—bots are being used to lubricate and augment influence activities, including facilitating content amplification and community cultivation efforts. They are standing in for official Islamic State operatives and advocates, connecting people with the movement based on common behaviors, shared interests, and/or ideological proximity while minimizing risk for the broader organization.

As we have shown above, the bots we identified during the data collection period are mainly clustered around two communities of Islamic State supporters, with one group seemingly focusing on facilitating discussion and exchange, and another augmenting content distribution efforts. Crucially, as things stand, they are not "intelligent." Indeed, we did not identify any that use language models or intelligent systems to interact or generate content, even though such a capability is technically possible. Instead, these bots deploy basic scripts to execute specific—and simple—tasks. Notwithstanding their simplicity, though, they are clearly and fundamentally embedded within the Islamic State's online ecosystem. Together, they directly facilitate its survival, amplifying and minimizing risk when it comes to the spread of content and, more importantly, keeping the consumers of this content within reach of it—something that has become increasingly difficult to do in recent years due to Telegram's highly aggressive and well-targeted disruption efforts.

Understanding how these bot communities operate and, more to the point, on what basis they are structured is critical if their efforts are to be meaningfully and sustainability undermined any time soon. On that basis, future research would do well to build on these findings, perhaps studying temporal aspects of the network with a view to tracing its evolution and robustness. Another option for further study would be to address one of the major limitations of this study: the fact that that, due to ethical considerations, content posted by human users is missing from the analysis, and only content posted by the bots as out-of-context conversations is explored. Analyzing how bot content overlaps and interacts with human user content is a logical and necessary next step for exploration.

Acknowledgments

We thank Felipe Vaca for his help in reviewing some paragraphs. The research reported in this work was partially supported by the EU H2020 ICT48 project "Humane AI Net" under contract # 952026 and by the European Union—Horizon 2020 Program under the scheme "INFRAIA-01-2018-2019"—Integrating Activities for Advanced Communities," Grant Agreement n.871042, "SoBigData++: European Integrated Infrastructure for Social Mining and Big Data Analytics." Abdullah Alrhmoun also thanks NSF Grant Accelnet "Multinet" Grant # 1927425 for support.

Disclosure statement

No potential conflict of interest was reported by the authors.

Funding

This work was partially supported by the European Commission and the AccelNet-MultiNet program.

Notes on contributors

Abdullah Alrhmoun is a PhD candidate in network science at the Department of Network and Data Science at Central European University and holds a master's degree in engineering. His work revolves around the real-world applications of data analytics, including understanding complex systems through big data, text mining of extremist content, and developing tools to detect fake imagery. His research also explores extremism, disinformation and misinformation diffusion and the viral communication of online hate. He is also a visiting scholar at the Observatory on Social Media at Indiana University, researching the role of social bots in online extremist communities.

Charlie Winter is Director of Research at the threat intelligence platform ExTrac AI. Over the last decade he has worked in a range of academic positions in the US and UK, researching how and why violent extremists use strategic communication to further their political and military agendas in both on- and off-line spaces. Charlie's research has been supported by the Global Internet Forum for Counter-Terrorism (GIFCT), Facebook, the UK Home Office, and the US Department for Homeland Security, among others. Besides this, Charlie is an Associate Fellow at the International Centre for Counter-Terrorism and an Associate of the Imperial War Museum Institute in London. He is also member of the RESOLVE Network Research Advisory Council.

János Kertész (Dr. rer. nat. Eötvös University, Budapest). After holding positions in Budapest, Cologne, and Munich he was professor at the Budapest University of Technology and Economics, where he was director of the Institute of Physics 2001-20012. Since 2012 he has been professor at the Central European University (CEU) where he is head of the Department of Network and Data Science. His main interest is in interdisciplinary applications of statistical physics, including econophysics, network science and computational social science. He is elected member of the Hungarian Academy of Sciences and external member of the Finnish Academy of Science and Letters.

ORCID

Abdullah Alrhmoun http://orcid.org/0000-0002-5600-6482
János Kertész http://orcid.org/0000-0003-4957-5406

Notes

1. Neil Johnson, "New Terrorism Reveals New Physics," *Dynamics* 103 (2009): 148701.
2. Paul Staniland, "Organizing Insurgency: Networks, Resources, and Rebellion in South Asia," *International Security* 37, no. 1 (2012): 142–77; Paul Stephen Staniland, "Explaining Cohesion, Fragmentation, and Control in Insurgent Groups" (PhD diss., Massachusetts Institute of Technology, 2010); Cynthia Stohl, and Michael Stohl, "Networks of Terror: Theoretical Assumptions and Pragmatic Consequences," *Communication Theory* 17, no. 2 (2007): 93–124; Jonathan David Farley, "Breaking Al Qaeda Cells: A Mathematical Analysis of Counterterrorism Operations (A Guide for Risk Assessment and Decision Making)," *Studies in Conflict & Terrorism* 26, no. 6 (2003): 399–411; Stuart Koschade, "A Social Network Analysis of Jemaah Islamiyah: The Applications to Counterterrorism and Intelligence," *Studies in Conflict & Terrorism* 29, no. 6 (2006): 559–75.
3. Albert-László Barabási, "Network Science," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 371, no. 1987 (2013): 20120375.
4. Marc Sageman, *Understanding Terror Networks* (University of Pennsylvania Press, 2004).
5. Pedro Manrique, Zhenfeng Cao, Andrew Gabriel, John Horgan, Paul Gill, Hong Qi, Elvira M. Restrepo, et al., "Women's Connectivity in Extreme Networks," *Science Advances* 2, no. 6 (2016): e1501742.
6. Matteo Gregori, and Ugo Merlone, "Comparing Operational Terrorist Networks," *Trends in Organized Crime* 23, no. 3 (2020): 263–88.
7. Yahui Tian, and Yulia R. Gel, "Fusing Data Depth with Complex Networks: Community Detection with Prior Information," *Computational Statistics & Data Analysis* 139 (2019): 99–116; Gian Maria Campedelli, Iain Cruickshank, and Kathleen M. Carley, "A Complex Networks Approach to Find Latent Clusters of Terrorist Groups," *Applied Network Science* 4, no. 1 (2019): 1–22.
8. Matthew C. Benigni, Kenneth Joseph, and Kathleen M. Carley, "Online Extremism and the Communities that Sustain it: Detecting the ISIS Supporting Community on Twitter," *PLoS One* 12, no. 12 (2017): e0181405.
9. Bennett Clifford, and Helen Powell, "Encrypted Extremism: Inside the English-Speaking Islamic State Ecosystem on Telegram," *The George Washington University Program on Extremism* (2019): 27–53; Ahmad Shehabat,

Teodor Mitew, and Yahia Alzoubi, "Encrypted Jihad: Investigating the Role of Telegram App in lone Wolf Attacks in the West," *Journal of Strategic Security* 10, no. 3 (2017): 27–53.
10. Darren Loucaides, "How Telegram Became the Anti-Facebook," *Wired. Conde Nast*, February 8, 2022, https://www.wired.com/story/how-telegram-became-anti-facebook/.
11. Emilio Ferrara, Onur Varol, Clayton Davis, Filippo Menczer, and Alessandro Flammini, "The Rise of Social Bots," *Communications of the ACM* 59, no. 7 (2016): 96–104.
12. Felix Brünker, Julian Marx, Björn Ross, Stefan Stieglitz, and Milad Mirbabaie, "'The Tireless Selling-Machine'-Commercial Deployment of Social Bots during Black Friday Season on Twitter," in *15th International Conference on Wirtschaftsinformatik (Zentrale Tracks)* (2020), 1522–27.
13. Chengcheng Shao, Giovanni Luca Ciampaglia, Onur Varol, Alessandro Flammini, and Filippo Menczer, "The Spread of Fake News by Social Bots," *arXiv preprint arXiv:1707.07592* 96 (2017): 104.
14. Chengcheng Shao, Giovanni Luca Ciampaglia, Onur Varol, Kai-Cheng Yang, Alessandro Flammini, and Filippo Menczer, "The Spread of Low-Credibility Content by Social Bots," *Nature Communications* 9, no. 1 (2018): 1–9.
15. Wen Chen, Diogo Pacheco, Kai-Cheng Yang, and Filippo Menczer, "Neutral Bots Reveal Political Bias on Social Media," *arXiv preprint arXiv:2005.08141* (2020).
16. Alessandro Bessi, and Emilio Ferrara, "Social Bots Distort the 2016 US Presidential Election Online Discussion," *First Monday* 21, no. 11–7 (2016): 1–14.
17. Samuel Woolley, and Mark Kumleben, "Social Bots for Peace: Combating Automated Control with Automated Civic Engagement?" in *Social Media as a Space for Peace Education* (Cham: Springer, 2020), 203–223.
18. "Telegram FAQ, "Telegram," https://telegram.org/faq?setln=en (accessed October 7, 2022).
19. Gabriel Weimann, *Terrorism in Cyberspace: The Next Generation* (Columbia University Press, 2015); Maura Conway, Lee Jarvis, and Orla Lehane, eds. *Terrorists' Use of the Internet: Assessment and Response*, Vol. 136 (Ios Press, 2017).
20. Bruce Hoffman, *Inside Terrorism* (Columbia University Press, 1998).
21. Manuel Ricardo Torres-Soriano, "The Dynamics of the Creation, Evolution, and Disappearance of Terrorist Internet Forums," *International Journal of Conflict and Violence (IJCV)* 7, no. 1 (2013): 164–78.
22. Andrew Glazzard, "ISIS: The State of Terror," (2015), 85–87.
23. Amarnath Amarasingam, Shiraz Maher, and Charlie Winter, "How Telegram Disruption Impacts Jihadist Platform Migration," *Centre for Research and Evidence on Security Threats* (2021).
24. Aaron Brantly, "Banning Encryption to Stop Terrorists: A Worse than Futile Exercise," *CTC Sentinel* 10, no. 7 (2017): 29–33.
25. Charlie Winter, Abdullah Alrhmoun, and Abdul Sayed, "The Taliban's Vast Propaganda Machine has a New Target," *WIRED UK*, August, 2021, https://www.wired.co.uk/article/taliban-propaganda-news-afghanistan.
26. Joas Wagemakers, "Al-Qa 'ida's Editor: Abu Jandal al-Azdi's Online Jihadi Activism," *Politics, Religion & Ideology* 12, no. 4 (2011): 355–69; Christopher Anzalone, "Continuity and Change: The Evolution and Resilience of Al-Shabab's Media Insurgency, 2006–2016," *Hate Speech International* 9 (2016): 1–40.
27. Lina Khatib, *Image Politics in the Middle East: The Role of the Visual in Political Struggle* (Bloomsbury Publishing, 2012); Charlie Winter and Abdullah Alrhmoun, "Mapping The Extremist Narrative Landscape in Afghanistan," 2021, https://public-assets.extrac.io/reports/ExTrac_Afghanistan_201120.pdf.
28. ISIS Watch, "Telegram," https://t.me/ISISwatch (accessed October 7, 2022).
29. J. M. Berger, "How Terrorists Recruit Online (And How to Stop It)," *Brookings Institution* 9 (2015); Jytte Klausen, "Tweeting the Jihad: Social Media Networks of Western Foreign Fighters in Syria and Iraq," *Studies in Conflict & Terrorism* 38, no. 1 (2015): 1–22; Jonathon M. Berger, and Jonathon Morgan, "The ISIS Twitter Census: Defining and Describing the Population of ISIS Supporters on Twitter," (2015).
30. Bruce Hoffman, "Using the Web as a Weapon: The Internet as a Tool for Violent Radicalization and Homegrown Terrorism," *Statement Provided to the US House Committee on Homeland Security (November 6, 2007)* (2009); Bruce Hoffman, "The Myth of Grass-Roots Terrorism-Why Osama bin Laden Still Matters," *Foreign Affairs* 87 (2008): 133.
31. Maura Conway, "From al-Zarqawi to al-Awlaki: The Emergence and Development of an Online Radical Milieu," *CTX: Combating Terrorism Exchange* 2, no. 4 (2012): 12–22.
32. Bennett Clifford, "Migration Moments: Extremist Adoption of Text-Based Instant Messaging Applications," *Retrieved from London* (2020); Clifford and Powell, "Encrypted Extremism,"; Amarasingam et al., "How Telegram Disruption Impacts Jihadist Platform Migration."
33. J. Berger, "How ISIS Games Twitter," June 2014, https://www.theatlantic.com/international/archive/2014/06/isis-Iraq-twitter-social-media-strategy/372856/.
34. Oz Sultan, "Tackling Disinformation, Online Terrorism, and Cyber Risks into the 2020s," *The Cyber Defense Review* 4, no. 1 (2019): 43–60; Matthew Bondy, "Bad Bots." *The Project on International Peace and Security, Institute for the Theory and Practice of International Relations, College of William and Mary* (2017): 2016–17.
35. Emmanuel C. Ogu, Michael I. Ogu, and Chiemela Ogu, "Insights from Terrorism Intelligence and Eradication Efforts-Al-Qaeda, ISIS, Boko Haram-for more Pragmatic botnet Countermeasures," *International Journal of Collaborative Intelligence* 1, no. 4 (2016): 258–74.
36. Shehabat et al., "Encrypted Jihad."
37. Telegram, "Bots: An Introduction for Developers," https://core.telegram.org/bots.

38. Ibrahim B. Syed, "Birds in the Quran: The Hoopoe," 2021, https://aboutislam.net/muslim-issues/science-muslim-issues/birds-quran-hoopoe/.
39. Telegram, "Focused Privacy, Discussion Groups, Seamless Web Bots and More," https://telegram.org/blog/privacy-discussions-web-bots#broadcasts-meet-group-chats.
40. Santo Fortunato, and Darko Hric, "Community Detection in Networks: A User Guide," *Physics Reports* 659 (2016): 1–44.
41. To reiterate, we define the word "community" as it is customarily understood in network science—i.e., as a topological entity based on the connectedness of its constituent nodes; Andrea Lancichinetti, Santo Fortunato, and János Kertész, "Detecting the Overlapping and Hierarchical Community Structure in Complex Networks," *New Journal of Physics* 11, no. 3 (2009): 033015.
42. Michelle Girvan, and Mark E. J. Newman, "Community Structure in Social and Biological Networks," *Proceedings of the National Academy of Sciences* 99, no. 12 (2002): 7821–26.
43. M. Puck Rombach, Mason A. Porter, James H. Fowler, and Peter J. Mucha, "Core-Periphery Structure in Networks," *SIAM Journal on Applied Mathematics* 74, no. 1 (2014): 167–90.
44. Stephen P. Borgatti, and Martin G. Everett, "Models of Core/Periphery Structures," *Social Networks* 21, no. 4 (2000): 375–95.