

Computational Social Science to Gauge Online Extremism

Emilio FERRARA ^{a,1}

^a*University of Southern California, Information Sciences Institute, USA*

Abstract. Recent terrorist attacks carried out on behalf of ISIS on American and European soil by lone wolf attackers or sleeper cells remind us of the importance of understanding the dynamics of radicalization mediated by social media communication channels. In this paper, we shed light on the social media activity of a group of twenty-five thousand users whose association with ISIS online radical propaganda has been manually verified. By using a computational tool known as *dynamical activity-connectivity maps*, based on network and temporal activity patterns, we investigate the dynamics of social influence within ISIS supporters. We finally quantify the effectiveness of ISIS propaganda by determining the adoption of extremist content in the general population and draw a parallel between radical propaganda and epidemics spreading, highlighting that information broadcasters and influential ISIS supporters generate highly-infectious cascades of information contagion. Our findings will help generate effective countermeasures to combat the group and other forms of online extremism.

Keywords. computational social sciences, social media, extremism, social influence

1. Introduction

Researchers in the computational social science community have recently demonstrated the importance of studying online social networks to understand our society [33]. New powerful technologies are usually harbinger of abuse, and online platforms are no exception: social media have been shown to be systematically abused for nefarious purposes [19]. As online social environments yield plenty of incentives and opportunities for unprecedented, even “creative” forms of misuse, single individuals as well as organizations and governments have systematically interfered with these platforms, oftentimes driven by some hidden agenda, in a variety of reported cases:

- During crises, social media have been effectively used for emergency response; but fear-mongering actions have also triggered mass hysteria and panic [26,19].
- Political conversation has been manipulated by means of orchestrated astroturf campaigns [39,36] even during election times [27,11].
- Anti-vaccination movements [47,49], as well as conspiracy (and other anti-science) theorists [10,18], took social media by the storm and became responsible for a major health crisis in the United States [42].

¹Corresponding Author: Emilio Ferrara, USC Information Sciences Institute, 4676 Admiralty Way #1001, Marina del Rey, CA (USA) 90292; E-mail: emiliofe@usc.edu

January 2017

- Social media bots (non-human automated accounts) have been used to coordinate attacks to successfully manipulate the stock market, causing losses in the billions of dollars [28,19,20].
- Some governments and non-state actors have been active on social media to spread their propaganda. In some cases, they have allegedly “polluted” these platforms with content to sway public opinion [19,30], or to hinder the ability of social collectives to communicate, coordinate, and mobilize [46].

Especially related to the last point, researchers have been recently devoting more attention to issues related to online propaganda campaigns [44,40,5]. Increasing evidence provided by numerous independent studies suggests that social media played a pivotal role in the rise in popularity of the *Islamic State of Iraq and al-Sham* (viz. ISIS) [23,45, 53]. Determining whether ISIS benefitted from using social media for propaganda and recruitment was central for many research endeavors [8,6,35].

Analyses by Berger and Morgan suggested that a restricted number of highly-active accounts (500-1000 users) is responsible for most of ISIS’ visibility on Twitter [6]. However, Berger’s subsequent work suggested that ISIS’ reach (at least among English speakers) has stalled for months as of the beginning of 2016, due to more aggressive account suspension policies enacted by Twitter [7]. Other researchers tried to unveil the roots of support for ISIS, suggesting that ISIS backers discussed Arab Spring uprisings on Twitter significantly more than users who stood against ISIS [35].

These early investigations all share one common methodological limitation: borrowing the data collection approach from previous research [25,15,51], these studies start from keywords known to be associated to ISIS to collect social network data [6,7,35]. We argue that, unfortunately, it is not sufficient to focus on keyword-based online chatter to pinpoint to relevant actors of radical conversation, and in particular to identify those who contribute generating and spreading the radical message, as opposed to those who consume it. Furthermore, keyword-based data collection practices introduce biases and noise by capturing users who casually adopt the pre-selected keywords, or who share news that happen to include those keywords but do not represent propaganda efforts. One contribution of our work is to address this issue by focusing on manually-verified set of ISIS supporter accounts. More in general, this study aims to address the following first research question:

RQ1: Can we define a solid methodological framework and suggest good practices for data collection, validation, and analysis to study online radicalization?.

Setting such best practices will hopefully steer the computational social sciences research community in the direction of producing more rigorous and reproducible work.

After describing how we collected and curated the dataset object of this study, we move forward to investigate the dynamics of online propaganda adopted by ISIS. Using computational social science tools to gauge online extremism, we aim to answer the following second research question:

RQ2: What types of network and temporal patterns of activity reflect the dynamics of social influence within ISIS supporters? And, can we quantify the adoption of extremist content in the general population?.

Our findings will help generate effective countermeasures to combat the group and other forms of online extremism.

January 2017

2. Data Collection and Curation

Due to the the limits of keyword-based data collection approaches we highlighted above, in this work we exclusively rely on data and labels obtained by using a procedure of manual data curation and expert verification. We obtained a list of Twitter accounts whose activity was labeled as supportive of the Islamic State by the online crowd-sourcing initiative called *Lucky Troll Club*. The goal of this initiative was to leverage annotators with expertise in Arabic languages to identify accounts affiliated to ISIS and report them to Twitter to request their suspension. Twitter’s anti-abuse team manually verifies all suspension requests, and grants some based on evidence of violation of Twitter’s Terms of Service policy against the usage of the platform for extremist purposes.

We verified that 25,538 accounts present in the *Lucky Troll Club* list have been suspended by Twitter in the period between March 17, 2015 and June 9, 2015. In this study, we focus only on this subset of twenty-five thousand accounts: we consider their activity to be unequivocally linked to the Islamic State, as determined by the two-step manual verification process described above. This list is available upon request from the author. For each account, we have at our disposal information about the suspension date, as well as the number of followers of that user as of the suspension date.

2.1. Twitter data collection

The next step of our study consisted in collecting data related to the activity of the 25,538 ISIS supporters on Twitter. To this purpose, we leveraged the Observatory on Social Media (OSoMe) data source set up by our collaborators at Indiana University [17], which continuously collects the Twitter data stream from the *gardenhose* API (roughly a 10% random sample of the full Twitter data stream). Using this large data stream avoids known issues derived by using the public Twitter stream API which serves only less than 1% of the overall tweets [37].

We obtained all tweets present in the OSoMe database that have been posted by any of the twenty-five thousand ISIS accounts prior to their suspension. We also included all retweets these tweets generated, and all tweets containing mentions to such set of ISIS supporters. The resulting dataset that we will study consists of 3,395,901 tweets. Almost 1.2 million of these tweets was generated by the ISIS accounts during the period between January 2014 and June 2015. We found that a total of 54,358 distinct other users has retweeted at least once one of the twenty-five thousand ISIS supporters. This amounts for the remainder of about 2.2 million tweets in our dataset.

Summarizing, we identified two best practices to answer **RQ1**:

- Favor starting from a manually-verified list of users involved in online propaganda, radicalization efforts, or recruitment, rather focusing on keyword-based searches. When such manually validated lists are not available, human annotations can be generated by means of services such as Amazon Mechanical Turk.
- Use large social media data streams, when available, rather than small samples that can be biased. Services like the Indiana University’s OSoMe database [17] can provide data sources especially valuable for Twitter-based social media studies. Alternatively, the Twitter Search API can yield comprehensive data around specific users or topics, provided that the search is limited to short time frames.

January 2017




Figure 1. Timeline of suspensions/day of the 25,538 ISIS-supporting accounts. The main figure adopts a logarithmic scale (y axis) while the inset figure shows the original data.

3. Results

We next report our investigation to address **RQ2**. We study the activity of ISIS supporters and sympathizers on Twitter by means of the data collected as described above.

3.1. Activity and Support of ISIS Accounts

The first form of validation that we performed pertains the mechanisms of suspension of ISIS accounts on Twitter. As we mentioned before, we identified a set of twenty-five thousand accounts related to ISIS that have been suspended: Fig. 1 shows the timeline of suspensions of the accounts under investigation. The suspension period occurred throughout almost three months, with the first suspensions occurring on March 17, 2015 and the last occurring on June 9, 2015. After this date, none among the twenty-five thousand accounts in our list is anymore active on Twitter. Account suspensions appear to occur in batches, some more substantial than others, with a significant spike of suspensions (over ten thousands) occurred on April 2, 2015. Our findings are consistent with *The Guardian*'s report that, between April 2015 and February 2016, Twitter's anti-abuse task force suspended more than 125,000 accounts linked to ISIS [54].

Recent literature reported contrasting evidence about the activity and popularity of ISIS supporters on Twitter [6,7]. As for today, it is unclear whether ISIS accounts obtained a significant support on Twitter, and to what extent they were active. Berger and collaborators first found that ISIS presence was very pervasive on Twitter during 2014-2015 [6], and later suggested that only a core of 500-2000 ISIS users was active after that period [7]. To shed light on this question, we calculated the distribution of the number of followers and followees (friends) of ISIS accounts at the time of their suspension, along with the number of tweets and retweets they generated. These four distributions are shown in Fig. 2. Let us discuss support and activity separately.

January 2017




Figure 2. Distributions of the number of tweets posted by ISIS supporters, retweets generated, and number of followers and followees at the time of their suspension.

Concerning ISIS support on Twitter, we notice that the distribution of followers of ISIS accounts exhibits the long tail typical of Twitter [32] and other social networks (*cf.* yellow dash-dotted line in Fig. 2). This skewed distribution has mean $\mu = 516$ ($\sigma = 1,727$), median $Q_2 = 130$, and lower and upper quartiles $Q_1 = 37$ and $Q_3 = 401$. This means that the majority of accounts has a limited number of followers (for example, one quarter of the users has less than 37 followers), yet a significant number of ISIS supporters managed to obtain a large number of followers before getting suspended (in fact, the upper quartile of users has more than 400 followers). The presence of this broad distribution of followership suggests that influence and radicalization operations of Islamic State supporters on Twitter were successful for at least several thousand of their accounts. This is in line with Berger’s early results discussed above [6]. Another interesting insight is yielded by the distribution of followees (*cf.* green dotted line in Fig. 2): differently from the distribution of followers, this distribution shows an unexpected upward trend in the regime between 100 and 1,000 followees. This characteristic behavior has been associated to forms of social network manipulation, for example attempts to create rings or cliques in which multiple accounts under the control of a same entity all follow each other to reciprocally increase their visibility and followership [39]. We thus suggest that ISIS accounts enacted strategies to artificially enhance their visibility by strengthening one another social networks.

A significant portion of ISIS accounts was very active on Twitter: both distributions of tweets (*cf.* solid blue line) and retweets (*cf.* red dashed line) shown in Fig. 2 exhibit the typical power-law shape common to social networks with heterogeneous activity patterns. This implies that a significant fraction of users posted and retweeted large amounts of tweets. For example, at least 1% of the ISIS users posted at least 30 tweets during the observation period; similar figures hold for retweeting. More importantly, there appears to be a strong core constituted by a few dozen accounts who posted and retweeted hundreds of tweets in the same period. This set of very active ISIS supporters we found is compatible with what reported by Berger’s first study [6]. Furthermore, there seems to

January 2017




Figure 3. Timeline of total number of tweets generated by ISIS supporters. The main figure shows the number of tweets per week; the inset figure shows the cumulative count.

be a handful of accounts with thousands of tweets and retweets, suggesting the likely presence of some social media bot used to enhance the volume of content generated by ISIS and its spreading on Twitter [20].

To investigate ISIS accounts' activity further, we extrapolated the time series of the total volume of tweets and retweets posted every week by the ISIS users under investigation: the result is shown in Fig. 3. It's worth noting that our observation window spans 1.5 years and starts on January 2014 when the firsts among the twenty-five thousand ISIS supporters became active on Twitter: Although the activity volume slowly builds up over time, it is only in early March 2015 (15 months into our observation period) that the volume of tweets per week drastically spikes.

In the early regime between January 2014 and March 2015 the volume of tweets associated to ISIS accounts spans 1,000 to 10,000 per week. This increases nearly tenfold after March 2015, with a spike of over 100,000 tweets per week, and an average of over 60,000 tweets per week in the period between March and May 2015. This period concurs with the period of strongest Twitter suspensions shown in Fig. 1, suggesting a timely reaction of Twitter to fight the activity of ISIS users on the platform. Indeed, the volume of tweets per week produced by these accounts drops in early June, and goes to zero, as expected, in the late period of observation when fewer and fewer of the ISIS accounts under investigation are left unchecked on Twitter. Cumulatively, the ISIS users under investigation produced almost 1.2 million tweets during the observation period (*cf.* inset of Fig. 3).

Our findings pinpoint to the power of the crowd-sourcing volunteer initiative that set to bring up to Twitter's attention these accounts. However, there is no evidence to quantify how many (if any) of the ISIS supporters not recorded by the *Lucky Troll Club* operation were independently suspended by Twitter anti-abuse team.




Figure 4. Dynamical Activity-Connectivity map.

3.2. Dynamical Activity-Connectivity Map

We conclude this analysis by exploring the relationship between activity and support. To investigate what effect the progression of activity levels of an ISIS user has of their connectivity evolution (and viceversa), in Fig. 4 we show a *Dynamical Activity-Connectivity map*. This map is a dynamic variant of the map proposed by Gonzalez-Bailon and collaborators—see Figure 4 in *Broadcasters and Hidden Influentials in Online Protest Diffusion* [24]. Fig. 4 shows the probability density of users in the two-dimensional space where the x-axis represents the growth of connectivity formation, and the y-axis conveys the rate of messaging activity. For a given user u , x_u and y_u are defined as

$$x_u = \frac{1 + \delta f_u}{1 + \delta F_u} \quad \text{and} \quad y_u = \frac{1 + m_u}{1 + M_u}.$$

We use the notation f_u and F_u to identify the number of followers and friends, respectively, of a user u . The variation of followers and friends of user u over a period of time t are thus defined as $\delta f = \frac{f_u^{\max} - f_u^{\min}}{t}$ and $\delta F = \frac{F_u^{\max} - F_u^{\min}}{t}$; the length of time t is defined as the number of days of u 's activity, measured from registration to suspension (this varies from user to user). Finally, m_u is the number of mentions user u received by others, and M_u is the number of mentions user u made to others, during u 's activity period. All values are added to the unit to avoid zero-divisions and allow for logarithmic scaling (i.e., in those cases where the variation is zero). The “heat” (the color intensity) in the map represents the joint probability density $pdf(x,y)$ for users with given values of x and y . The plot also introduce a bin normalization to account for the logarithmic binning.

The *Dynamical Activity-Connectivity map* we conceived is interpreted as follows: the bulk of the joint probability density mass should be observed in the neighborhood of $(1,1)$, that hosts the majority of accounts for which the variation of the two dimensions is comparable.

Let us discuss the two dimensions of *connectivity growth* and *activity rate* separately. The *connectivity growth* is captured by the x axis and, in our case, ranges roughly between 10^{-2} and 10^3 . Users for which $x > 1$ (i.e., 10^0) are those with a followership that grows much faster than the rate at which these users are following others. In other words, they are acquiring social network popularity (followers) at a fast-paced rate. Note that, if a user is acquiring many followers quickly, but s/he is also following many users at a similar rate, the value of x will be near 1. This is a good property of our measure because it is common strategy on social media platforms, especially among bots [20,11], to indiscriminately follow others in order to seek for reciprocal followerships. Our Dynamical Activity-Connectivity map will discriminate users with fast-growing followerships, who will appear in the right-hand side of the map, from those who adopt that type of reciprocity-seeking strategy. The former group can be associated with highly popular users with a fast-paced followership growth. According to Gonzalez-Bailon and collaborators [24] this category is composed by two groups: *influential users* and *information broadcasters*, depending on their activity rates. Values of $x < 1$ indicate users who follow others at a rate higher than that they are being followed; they fall in the left-hand side of the map. According to Gonzalez-Bailon and collaborators, these are mostly the *common users*, although the so-called *hidden influentials* also sit in this *low-connectivity* regime.

As for what concerns the y axis, it measures the *activity rate*, i.e., the rate at which a user receives mentions versus how frequently s/he mentions others. Users with values of $y > 1$ are those who receive systematically more mentions with respect to how frequently they mention others. This group of users can be referred to as *influentials*, i.e., those who are referred to significantly more frequently than others in the conversation; they fall in the upper region of the map, and according to Gonzalez-Bailon *et al.*, depending on their connectivity growth can be divide in influential ($x > 1$) and hidden influential ($x < 1$) users. Conversely, users with values of $y < 1$ are those who generate increasingly more mentions over time, either because they reply to many tweets, or because they address directly other users. This group generally represents the common-user behavior ($x < 1$), although information broadcasters ($x > 1$) also exhibit the same *low-activity* rate. These users fall in the lower region of the map.

Now that a reading of dynamical activity-connectivity maps has been provided, we can proceed with interpreting Fig. 4: the bottom-left quadrant reports the most common users, those with both activity and connectivity growth lesser than 1. Conversely, the upper-right quadrant reports users with the higher connectivity growth and activity rates. These are influential ISIS supporters who are very active in the discussion. We note how the connectivity growth dimension spans three orders of magnitude in the positive domain, while the activity rate dimension only spans two orders of magnitude. This means that some users' followerships grows tens of times faster than the rate at which they follow others; conversely, the rate of receiving mentions is only up to tens of times higher than that of mentioning of others. In the next section, we will devote special attention to these four different classes of users to determine what types of differences emerge in the ISIS social network.

We further address **RQ1** by offering *Dynamical Activity-Connectivity maps* as a powerful tool for computational social scientists to study influence and authority dynamics of online extremism in social networks with temporal activity patterns.

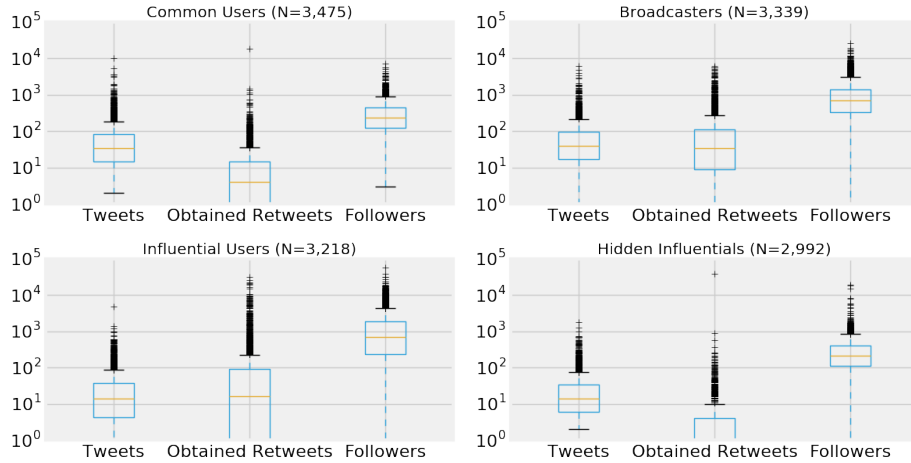


Figure 5. Distributions of (i) number of posted tweets, (ii) number of obtained retweets, and (iii) number of followers, for common users, broadcasters, influential, and hidden influential users, respectively.

3.3. Dynamical Classes of ISIS Supporter Behaviors

Prior research illuminated on the dynamical aspects of activity and connectivity in social media [34,24]. Next, we focus on the four classes of user behaviors highlighted by the dynamic activity-connectivity map. We first select, out of the twenty-five thousand ISIS supporters in our dataset, only the subset of those who have mentioned and have been mentioned at least once during the observation period. This will allow to focus on active accounts and correctly capture their activity rate. This filter reduces the number of users under investigation to $N = 13,024$ ISIS supporters, nearly half of the entire ISIS population initially collected. We further divide these users in the four dynamical classes defined above. The classification is obtained by simply adopting the rules defining the four quadrants of Fig. 4, which yields $N = 3,475$ common users ($x < 1, y < 1$), $N = 3,339$ information broadcasters ($x > 1, y < 1$), $N = 3,218$ influentials ($x > 1, y > 1$), and $N = 2,992$ hidden influential users ($x < 1, y > 1$).

For each of these users, we generated the distribution of (i) the total number of tweets they posted, (ii) the cumulative number of times they have been retweeted, and (iii) the maximum number of followers they gathered. Fig. 5 shows the boxplots corresponding to the four dynamical classes. Significant differences emerge: common users produce an amount of tweets very similar to that of broadcasters, but they accrue nearly one order of magnitude less retweets than the latter. Information broadcasters also appear to generate the largest followerships, on par with influential users; influentials, however, post significantly less tweets, while accumulating similar amounts of retweets than broadcasters, suggesting that our map successfully captures a notion of social influence intended as a proxy for attention generated to one's posts. The class of hidden influentials shows comparable activity to influential users, but significantly less influence, accruing about one order of magnitude less tweets and significantly smaller followerships than influentials.

Our analysis suggests that different classes of ISIS supporters' behaviors emerge. In the future, we will study what are the characteristics of different classes that produce the most effective propaganda and make the most influential users, analyzing content and language, political and religious beliefs, motives and attitudes of the ISIS social network.

January 2017




Figure 6. Timeline of the number of individuals per day who adopted ISIS contents.

3.4. Adoption of ISIS Propaganda

So far, our analysis focused on characterizing some dimensions of the behavior of ISIS supporters on Twitter. Next, we investigate whether the content they generated has been adopted by other users who have become exposed to it. Our notion of adoption is very simple: when a user who does not appear set of the twenty-five thousand ISIS accounts retweets for the first time any tweet generated by one such ISIS supporter, we count this as a content adoption. Although some recent work suggested that retweeting radical propaganda can be considered as an early sign of radicalization [6,35], we call for caution about interpreting the results in such a way: this definition greatly simplifies the notion and complexity of such types of online adoption processes [12,13,31,21,22]. However, we do believe that investigating the spread of radical content in the user population has a value to determine the extent and effectiveness of ISIS propaganda operations.

Fig. 6 shows the time series of the number of ISIS content adoptions per day during the first half of 2015. Prior to that, no significant amount of adoptions could be observed, partly due to the low activity rate of the accounts under investigation. In the three months between March and June 2015, we notice a significant uptake in the number of adoptions, peaking at nearly one thousand adoptions per day. In that period, at least 10,000 tweets per day (70,000-100,000 tweets/week) were generated by ISIS accounts (*cf.* Fig 3). This suggests that a very significant fraction of tweets, about 5-10%, was actually retweeted on average at least once by other users. During this period, a total of 54,358 distinct other users has retweeted at least once one of the twenty-five thousand ISIS supporters.

If we simplify ISIS propaganda diffusion as an infectious disease, we can draw a parallel with epidemics. The *basic reproduction number* R_0 of an infection is the number of cases generated on average by an infected individual over the course of its infectious period, in an otherwise uninfected population. Given that 25,538 ISIS supporters generated 54,358 distinct *infected users*, we can derive an $R_0 = 2.13$ for the ISIS propaganda “infection”. In other words, an ISIS supporter before being suspended on average “infected” 2.13 other users. For comparative purposes, health epidemics like Ebola, SARS, HIV/AIDS, and certain strains of influenza all have similar values of $2 < R_0 < 3$.




Figure 7. Distributions of retweets (top) and mentions (bottom) obtained by ISIS users as divided in the four dynamical classes, namely common users, broadcasters, influential and hidden influential users. Each group reports the size of the population under investigation (different in the two scenarios).

3.5. Contagion dynamics of ISIS Propaganda

We conclude our analysis by studying the cascades of content adoptions generated by the four classes of users defined above. This investigation is twofold: first, we would like to determine whether the mechanisms of receiving retweets (what we defined as content adoption) and being mentioned by out-of-sample users exposed to extremist content follow the same or different dynamics. Second, we will compute the distributions of scores for the basic reproduction number R_0 relative to both receiving retweets and mentions for the four classes of users. The goal is to reveal whether any significant difference emerge between groups of users in their content spreading efficacy, and ultimately to understand which groups of users generated the most effective information contagions.

Fig. 7 shows the distributions of retweets (top) and mentions (bottom) received by the users in the four classes defined by means of the dynamical activity-connectivity map. Although overall all distributions are broad, as expected given the heterogeneous nature of information diffusion, the dynamics of obtaining retweets and mentions are significantly different for the four groups: in particular, for what concerns receiving mentions, no appreciable difference emerges among the four classes of users, which suggests that ISIS supporters are being mentioned in a similar fashion regardless of the class they belong to. However, receiving retweets shows a different mechanism: influential and broadcaster users receive generate significantly larger retweet cascades much more frequently than common users and hidden influentials.

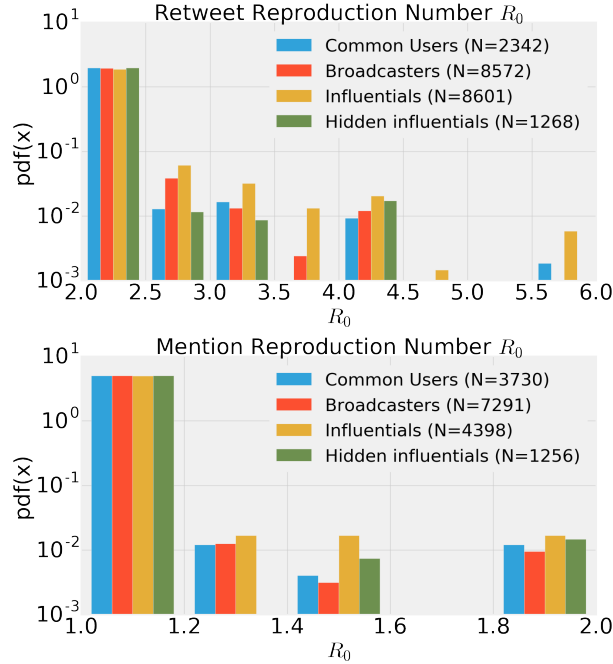


Figure 8. Probability density of basic reproduction scores R_0 relative to obtained retweets (top) and mentions (bottom) calculated for the four dynamical classes of ISIS users.

One question that thus rises concerns whether the actual contagion effectiveness varies between classes of users. In other words, what are the classes of users that are generating the most effective propaganda campaigns, in terms of adoption (*i.e.*, retweets) and engagement (*i.e.*, mentions)? To this purpose, for each ISIS user who obtained at least one retweet (resp., mention), we calculated the fraction $\frac{RT}{T}$ of received retweets RT (resp., mentions) over the total number of his/her tweets T that have been retweeted at least once (resp., mentioned) by an out-of-sample user (*i.e.*, a user not labeled as ISIS by our list). This measures the rate of diffusion of tweets in the otherwise uninfected population. We note that this is a simplification of the more traditional notion of information contagion where we would consider all the tweets generated and more complex diffusion mechanisms accounting e.g., for exposures, due to the limitation of the platform under study (namely, we do not have any information about information exposure on Twitter). With some abuse of notation, we thus consider the fraction $\frac{RT}{T}$ to convey the meaning of the basic reproduction number R_0 typical of epidemiology.

Finally, ISIS accounts are divided in their four classes, according to the usual classification, and the probability density of the basic reproduction scores R_0 is shown in Fig. 8, separately for the two dynamics of receiving retweets (top) and mentions (bottom). While no significant difference emerges, in either scenario, among the four different classes of users for the lowest scores (*i.e.*, $2 < R_0 < 2.5$ for retweets, and $1 < R_0 < 1.2$ for mentions), strong class differences emerge for users whose content are more contagious ($R_0 > 2.5$ for retweets, and $R_0 > 1.2$ for mentions): concerning retweets, influential users, followed by information broadcasters, are receiving systematically more attention than users in other classes; the class differences for mentions are less pronounced.

4. Related work

One of the first computational frameworks, proposed by Bermingham *et al.* [9] in 2009, combined social network analysis with sentiment detection tools to study the agenda of a radical YouTube group: the authors examined the topics discussed within the group and their polarity, to model individuals' behavior and spot signs of extremism and intolerance, seemingly more prominent among female users. The detection of extremist content (on the Web) was also the focus of a 2010 work by Qi *et al.* [38]. The authors applied hierarchical clustering to extremist Web pages to divide them into different pre-imposed categories (religious, anti immigration, etc.).

Scanlon and Gerber proposed the first method to detect cyber-recruitment efforts in 2014 [43]. They exploited data retrieved from the Dark Web Portal Project [14], a repository of posts compiled from 28 different online fora on extremist religious discussions (e.g., Jihadist) translated from Arabic to English. After annotating a sample of posts as recruitment efforts or not, the authors use Bayesian criteria and a set of textual features to classify the rest of the corpus, obtaining good accuracy, and highlighted the most predictive terms.

Along the same trend, Agarwal and Sureka proposed different machine learning strategies [1,48,3,4] aimed at detecting radicalization efforts, cyber recruitment, hate promotion, and extremist support in a variety of online platforms, including YouTube, Twitter and Tumblr. Their frameworks leverage features of contents and metadata, and combinations of crawling and unsupervised clustering methods, to study the online activity of Jihadist groups on the platforms mentioned above.

A few studies explored unconventional data sources: one interesting example is the work by Vergani and Bliuc [52] that uses sentiment analysis (Linguistic Inquiry and Word Count [50]) to investigate how language evolved across the first 11 Issues of Dabiq, the flagship ISIS propaganda magazine. Their analysis offers some insights about ISIS radicalization motives, emotions and concerns. For example, the authors found that ISIS has become increasingly concerned with females, reflecting their need to attract women to create their utopia society, not revolving around warriors but around families. ISIS also seems to have increased the use of internet jargon, possibly to connect with the identities of young individuals online.

Concluding, two very recent articles [41,29] explore the activity of ISIS on social media. The former [41] focuses on Twitter and aims at detecting users who exhibit signals of behavioral change in line with radicalization: the authors suggest that out of 154K users only about 700 show significant signs of possible radicalization, and that may be due to social homophily rather than the mere exposure to propaganda content. The latter study [29] explores a set of 196 pro-ISIS aggregates operating on VKontakte (the most popular Russian online social network) and involving about 100K users, to study the dynamics of survival of such groups online: the authors suggest that the development of large and potentially influential pro-ISIS groups can be hindered by targeting and shutting down smaller ones.

For additional pointers we refer the interested reader to two recent literature reviews on this topic [16,2].

January 2017

5. Conclusions

Since the appearance of the Islamic State (viz. ISIS), a consensus has emerged on the relationship between extremism and social media, namely that ISIS' success as a terrorist organization is due at least in part to its savvy use of social media. It is widely believed that ISIS has managed to increase its roster to tens of thousands of members by broadcasting its savage attacks over social media platforms such as Twitter, which helps radicalize and ultimately recruit fighters from around the world. Recent attacks on American and European soil demonstrate ISIS' potential to reach, organize, and mobilize lone wolves and sleeper terrorist cells among westerners. Many of these actors are known to have consumed radical material online and many have claimed to gravitate towards Islamic State because of it. This paper sheds light on the social media activity of a group of twenty-five thousand users associated with ISIS. These accounts have been manually identified, reported to Twitter for verification, and subsequently suspended due to their involvement with radical propaganda. This process yielded a human-curated dataset containing over three million tweets generated during a period of one and half year. We adopt content-agnostic statistical and network techniques to dissect these users' temporal activities and connectivity patterns. By building upon a computational tool named *dynamical activity-connectivity map*, we highlighted the dynamics of social influence within ISIS supporters. Our findings suggest complex strategies carried out by these users to manipulate and influence others: four dynamical classes of ISIS supporters emerged (common sympathizers, information broadcasters, influential and hidden influential users), each with distinct activity and connectivity patterns. We concluded by quantifying the extent to which ISIS support and extremist content are adopted in the general population: by drawing a parallel between propaganda and epidemics spreading, we determined that each ISIS supporter "infected" on average 2.13 other users before Twitter suspended his/her account, highlighting that receiving retweets and mentions follow different dynamics, and that broadcasters and influential users generate much more widespread contagions. Although we call for caution in the interpretation of these results, due to the great simplifications introduced by our framework, we believe that our findings will help design and implement effective countermeasures and responses to ISIS social media operations and other forms of online extremist propaganda.

Acknowledgments

The author is grateful to Max Abrahms (Northeastern University) for useful discussions, and to Alessandro Flammini and Onur Varol (Indiana University) for their support in collecting the Twitter dataset. This work has been partly funded by the Office of Naval Research (ONR), grant no. N15A-020-0053. This research is also based upon work supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA). The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, ONR, or the U.S. Government. The U.S. Government had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

References

- [1] S. Agarwal and A. Sureka. A focused crawler for mining hate and extremism promoting videos on youtube. In *Proc. 25th ACM conference on Hypertext and social media*, pages 294–296. ACM, 2014.
- [2] S. Agarwal and A. Sureka. Applying social media intelligence for predicting and identifying on-line radicalization and civil unrest oriented threats. *arXiv preprint arXiv:1511.06858*, 2015.
- [3] S. Agarwal and A. Sureka. Using KNN and SVM based one-class classifier for detecting online radicalization on Twitter. In *Distributed Computing and Internet Technology*, pages 431–442. 2015.
- [4] S. Agarwal and A. Sureka. Spider and the flies: Focused crawling on tumblr to detect hate promoting communities. *arXiv preprint arXiv:1603.09164*, 2016.
- [5] W. H. Allendorfer and S. C. Herring. ISIS vs. the US government: A war of online video propaganda. *First Monday*, 20(12), 2015.
- [6] J. Berger and J. Morgan. The ISIS Twitter Census: Defining and describing the population of ISIS supporters on Twitter. *The Brookings Project on US Relations with the Islamic World*, 3(20), 2015.
- [7] J. Berger and H. Perez. The Islamic States diminishing returns on Twitter. *GWU Program on extremism*, 02-16, 2016.
- [8] J. Berger and B. Strathearn. Who matters online: measuring influence, evaluating content and countering violent extremism in online social networks. *International Centre for the Study of Radicalisation*, 2013.
- [9] A. Bermingham, M. Conway, L. McInerney, N. O’Hare, and A. F. Smeaton. Combining social network analysis and sentiment analysis to explore the potential for online radicalisation. In *2009 International Conference on Advances in Social Network Analysis and Mining*, pages 231–236. IEEE, 2009.
- [10] A. Bessi, M. Coletto, G. A. Davidescu, A. Scala, G. Caldarelli, and W. Quattrociocchi. Science vs conspiracy: Collective narratives in the age of misinformation. *PLoS one*, 10(2):e0118093, 2015.
- [11] A. Bessi and E. Ferrara. Social bots distort the 2016 us presidential election online discussion. *First Monday*, 21(11), 2016.
- [12] D. Centola. The spread of behavior in an online social network experiment. *Science*, 329(5996), 2010.
- [13] D. Centola. An experimental study of homophily in the adoption of health behavior. *Science*, 334(6060):1269–1272, 2011.
- [14] H. Chen, W. Chung, J. Qin, E. Reid, M. Sageman, and G. Weimann. Uncovering the dark Web: A case study of Jihad on the Web. *Journal of the American Society for Information Science and Technology*, 59(8):1347–1359, 2008.
- [15] M. D. Conover, E. Ferrara, F. Menczer, and A. Flammini. The digital evolution of Occupy Wall Street. *PLoS one*, 8(5):e64679, 2013.
- [16] D. Correa and A. Sureka. Solutions to detect and analyze online radicalization: a survey. *arXiv preprint arXiv:1301.4916*, 2013.
- [17] C. A. Davis, G. L. Ciampaglia, L. M. Aiello, K. Chung, M. D. Conover, E. Ferrara, A. Flammini, G. C. Fox, X. Gao, B. Gonçalves, et al. Osome: The iuni observatory on social media. *PeerJ Computer Science*, 2:e87, 2016.
- [18] M. Del Vicario, A. Bessi, F. Zollo, F. Petroni, A. Scala, G. Caldarelli, H. E. Stanley, and W. Quattrociocchi. The spreading of misinformation online. *Proceedings of the National Academy of Sciences*, 113(3):554–559, 2016.
- [19] E. Ferrara. Manipulation and abuse on social media. *ACM SIGWEB Newsletter*, (Spring):4, 2015.
- [20] E. Ferrara, O. Varol, C. Davis, F. Menczer, and A. Flammini. The rise of social bots. *Commun. ACM*, 59(7):96–104, 2016.
- [21] E. Ferrara and Z. Yang. Measuring emotional contagion in social media. *PLoS One*, 10(11), 2015.
- [22] E. Ferrara and Z. Yang. Quantifying the effect of sentiment on information diffusion in social media. *PeerJ Computer Science*, 1:e26, 2015.
- [23] A. Fisher. How jihadist networks maintain a persistent online presence. *Perspectives on Terrorism*, 9(3), 2015.
- [24] S. González-Bailón, J. Borge-Holthoefer, and Y. Moreno. Broadcasters and hidden influentials in online protest diffusion. *American Behavioral Scientist*, page 0002764213479371, 2013.
- [25] S. González-Bailón, J. Borge-Holthoefer, A. Rivero, and Y. Moreno. The dynamics of protest recruitment through an online network. *Scientific reports*, 1, 2011.
- [26] A. Gupta, H. Lamba, and P. Kumaraguru. \$1.00 per rt# bostonmarathon# prayforboston: Analyzing fake content on twitter. In *eCrime Researchers Summit (eCRS)*, 2013, pages 1–12. IEEE, 2013.
- [27] P. N. Howard and B. Kollanyi. Bots, #strongerin, and #brexit: Computational propaganda during the uk-eu referendum. *Available at SSRN 2798311*, 2016.

January 2017

- [28] T. Hwang, I. Pearce, and M. Nanis. Socialbots: Voices from the fronts. *Interactions*, 19(2):38–45, 2012.
- [29] N. Johnson, M. Zheng, Y. Vorobyeva, A. Gabriel, H. Qi, N. Velasquez, P. Manrique, D. Johnson, E. Restrepo, C. Song, et al. New online ecology of adversarial aggregates: ISIS and beyond. *Science*, 352(6292):1459–1463, 2016.
- [30] G. King, J. Pan, and M. E. Roberts. How the chinese government fabricates social media posts for strategic distraction, not engaged argument. *American Political Science Review (in press)*. Copy at <http://j.mp/1Txxiz1>, 2, 2017.
- [31] A. D. Kramer, J. E. Guillory, and J. T. Hancock. Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences*, 111(24):8788–8790, 2014.
- [32] H. Kwak, C. Lee, H. Park, and S. Moon. What is twitter, a social network or a news media? In *Proceedings of the 19th international conference on World wide web*, pages 591–600. ACM, 2010.
- [33] D. Lazer, A. S. Pentland, L. Adamic, S. Aral, A. L. Barabasi, D. Brewer, N. Christakis, N. Contractor, J. Fowler, M. Gutmann, et al. Life in the network: the coming age of computational social science. *Science (New York, NY)*, 323(5915):721, 2009.
- [34] J. Lehmann, B. Gonçalves, J. J. Ramasco, and C. Cattuto. Dynamical classes of collective attention in Twitter. In *Proceedings of the 21st international conference on World Wide Web*, pages 251–260, 2012.
- [35] W. Magdy, K. Darwish, and I. Weber. #FailedRevolutions: Using Twitter to study the antecedents of ISIS support. *First Monday*, 21(2), 2016.
- [36] P. T. Metaxas and E. Mustafaraj. Social media and the elections. *Science*, 338(6106):472–473, 2012.
- [37] F. Morstatter, J. Pfeffer, H. Liu, and K. M. Carley. Is the sample good enough? comparing data from twitter’s streaming api with twitter’s firehose. In *7th Int AAAI Conf on Weblogs and Social Media*, 2013.
- [38] X. Qi, K. Christensen, R. Duval, E. Fuller, A. Spahiu, Q. Wu, and C.-Q. Zhang. A hierarchical algorithm for clustering extremist web pages. In *2010 International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pages 458–463. IEEE, 2010.
- [39] J. Ratkiewicz, M. Conover, M. Meiss, B. Gonçalves, A. Flammini, and F. Menczer. Detecting and tracking political abuse in social media. *ICWSM*, 11:297–304, 2011.
- [40] S. Reardon. Terrorism: science seeks roots of terror. *Nature*, 517(7535):420–421, 2015.
- [41] M. Rowe and H. Saif. Mining pro-ISIS radicalisation signals from social media users. In *Proceedings of the Tenth International AAAI Conference on Web and Social Media*, pages 329–338, 2016.
- [42] S. Salzberg. Forbes: Anti-vaccine movement causes worst measles epidemic in 20 years. <http://www.forbes.com/sites/stevensalzberg/2015/02/01/anti-vaccine-movement-causes-worst-measles-epidemic-in-20-years/>.
- [43] J. R. Scanlon and M. S. Gerber. Automatic detection of cyber-recruitment by violent extremists. *Security Informatics*, 3(1):1–10, 2014.
- [44] Q. Schiermeier. Terrorism: Terror prediction hits limits. *Nature*, 517(7535):419, 2015.
- [45] J. Stern and J. M. Berger. *ISIS: The state of terror*. HarperCollins, 2015.
- [46] P. Suárez-Serrato, M. E. Roberts, C. Davis, and F. Menczer.
- [47] V. Subrahmanian, A. Azaria, S. Durst, V. Kagan, A. Galstyan, K. Lerman, L. Zhu, E. Ferrara, A. Flammini, and F. Menczer. The darpa twitter bot challenge. *Computer*, 49(6):38–46, 2016.
- [48] A. Sureka and S. Agarwal. Learning to classify hate and extremism promoting tweets. In *Intelligence and Security Informatics Conference (ISISIC), 2014 IEEE Joint*, pages 320–320. IEEE, 2014.
- [49] T. R. Tangherlini, V. Roychowdhury, B. Glenn, C. M. Crespi, R. Bandari, A. Wadia, M. Falahi, E. Ebrahimzadeh, and R. Bastani. “Mommy Blogs” and the vaccination exemption narrative: Results from a machine-learning approach for story aggregation on parenting social media sites. *JMIR Public Health and Surveillance*, 2(2), 2016.
- [50] Y. R. Tausczik and J. W. Pennebaker. The psychological meaning of words: LIWC and computerized text analysis methods. *Journal of language and social psychology*, 29(1):24–54, 2010.
- [51] O. Varol, E. Ferrara, C. L. Ogan, F. Menczer, and A. Flammini. Evolution of online user behavior during a social upheaval. In *Proceedings of the 2014 ACM conference on Web science*, pages 81–90, 2014.
- [52] M. Vergani and A.-M. Bliuc. The evolution of the ISIS’ language: a quantitative analysis of the language of the first year of Dabiq magazine. *Security, Terrorism, and Society*, 1(2):7–20, 2015.
- [53] M. Weiss and H. Hassan. *ISIS: Inside the army of terror*. Simon and Schuster, 2015.
- [54] D. Yadron. The Guardian: Twitter deletes 125,000 ISIS accounts and expands anti-terror teams. <http://www.theguardian.com/technology/2016/feb/05/twitter-deletes-isis-accounts-terrorism-online>.