Physica A 519 (2019) 169–180

Contents lists available at ScienceDirect

Physica A

journal homepage: www.elsevier.com/locate/physa

Temporal evolution of online extremist support

Zhenfeng Cao [a], Minzhang Zheng [b], Pedro D. Manrique [b], Zhou He [c], Neil F. Johnson [d,*]

[a] Department of Computing, The Hong Kong Polytechnic University, Hung Hom, Hong Kong
[b] Physics Department, University of Miami, Coral Gables, FL 33146, USA
[c] School of Economics and Management, University of Chinese Academy of Sciences, Beijing 100190, China
[d] Physics Department, George Washington University, Washington D.C. 20052, USA

HIGHLIGHTS

- We analyze social media concerning online support of violent extremism.
- Preferential attachment falls short of explaining the data.
- We propose a new model that explains the data by incorporating heterogeneity and network locality.
- Our analysis deepens understanding of clandestine or illicit online activities.

ARTICLE INFO

Article history:
Received 27 May 2018
Received in revised form 27 September 2018
Available online 26 December 2018

Keywords:
Online
Collective dynamics
Groups
Extremism

ABSTRACT

There is a significant amount of online human activity which is either clandestine or illicit in nature, and hence where individuals operate under fear of exposure or capture. Yet there is little theoretical understanding of what models best describe the resulting dynamics. Here we attempt to address this gap, by analyzing the evolutionary dynamics of the supporters behind the 95 pro-ISIS online communities (i.e. self-organized social media groups) that appeared recently on a global social media site. We show that their dynamical evolution can be explained by a model that incorporates effects of heterogeneity and network locality. Our analysis contributes to the understanding of online extremist support, and may also shed light on a broader spectrum of online human activities which are either clandestine or illicit in nature, and hence where individuals also operate under fear of exposure or capture.

1. Introduction

Just as the Internet can be used for good, it also serves as an ideal vehicle for the more clandestine and illicit side of human activity. For example, extremist entities such as ISIS (known as Islamic State) make ample use of the Internet for spreading their message and propaganda materials, recruiting young people, and soliciting funds. One particular social media platform VKontakte (www.vk.com), which has $\sim$ 350 million global users and is almost identical to Facebook, became the primary online social media source for ISIS propaganda and recruiting [1]. Unlike on Facebook where pro-ISIS activity is almost immediately eliminated, support on VKontakte develops around online groups (i.e. self-organized communities) which are akin to Facebook groups that support particular everyday topics such as a sports team. These online pro-ISIS groups may not only organize premeditated attacks, but also may help incite decentralized lone-wolf attacks [1]. Hence there is a compelling

* Correspondence to: Physics Department, Corcoran Hall, 725 21st Street NW, Washington, D.C. 20052, USA.
E-mail address: neiljohnson@gwu.edu (N.F. Johnson).

https://doi.org/10.1016/j.physa.2018.12.033
0378-4371/

need to investigate the dynamics of such online groups, especially during earlier stages of growth when group sizes are relatively small.

Here we analyze the evolutionary dynamics of the supporters behind the 95 pro-ISIS online communities (i.e. self-organized social media groups) that appeared on a global social media site in 2015 during the height of ISIS's online activity. Though focussed on ISIS support, our model, analysis and findings could serve in helping understand a broad spectrum of online human activities which are either clandestine or illicit in nature, and hence where individuals operate under fear of exposure or capture. Indeed, outside of extremism and online illicit or clandestine activity there is already increasing interest in understanding how communities, users, or groups attract new followers/members, and develop over time [2–10]. Many previous studies of grouping effects highlight the role played by the well-known preferential attachment (PA) effect in the formation of groups or equivalent network clusters [2,3,5,11–13]. We find that conventional PA model mechanisms, which by construction require non-local information about the network, cannot explain the unusual rapid and heterogeneous growths at the early growing stage of the online pro-ISIS groups observed in the empirical data (Section 2). We believe that the shortcoming of such preferential attachment mechanisms for online extremist support is possibly due to one or more of the following reasons:

1. The online pro-ISIS supporters and the self-organized groups that they form online, are under pressure. The members in such a group are discussing an extreme topic and supporting an illegal terrorist organization. These individuals have to co-habit the same online space as opposing individuals and entities (e.g. the online organization called Anonymous) as well as government agencies, all of whom are not only trying to defuse the narrative of the extremist social media groups but are also possibly trying to track down the identities of particular group members. As a result, the extremist supporters are under continual pressure and likely act online in ways that help them maintain a more hidden profile. This challenges significantly the typical PA model assumption of implicit knowledge of all group sizes across the whole population.
2. Extremism discussed in online pro-ISIS groups is a specialist niche topic. It seems less likely that people would be drawn to it simply because it is popular among others. Therefore, the probability of attaching should not only depend on node connectivity, but also should incorporate individual heterogeneity and network locality — just as readers of an article in a physics journal are likely to be physicists, and within that subpopulation, readers of an article on network science are more likely to be from the network science community. These are often overlooked factors that can play an important role in group formation and dynamics [14].
3. The evolution of online groups is affected by moderators who have the right to ban a group. Conventional PA models do not consider the consequences of such moderators. Their special role and powers suggest that their presence and activity can influence heavily the group evolution [15]. While a full theory that includes them would require multi-species analysis that we cannot yet provide, it suffices to say that their presence is another reason that the system dynamics can move beyond PA.

These issues need to be addressed in order to get a more complete picture of human activity online.

In the following sections of this paper, we go partway toward trying to address these issues. In so doing, our work also goes partway toward addressing the following wish-list in the network literature: (1) Many previous studies are based on particular definitions and techniques for detecting the groups. It would be very useful to have a more general framework for discussing what a group is [4–6,16]. (2) Although many previous studies focus either on the observational aspects [2,3,8–10] or theoretical modeling [6,17,18], knowledge about how network theories agree with empirical observations on both the microscopic and the macroscopic scales is rare. For instance, a large portion of studies focus on the global statistical properties such as the scaling behavior of the group/cluster size distribution and the evolution of the globally averaged quantities [2,3,5,11], yet little is known about the evolution at the group-specific level. This includes the study in Ref. [1] which focuses on the mechanisms producing the tail in the size distribution at larger group sizes. More detailed group-specific studies for any group size are needed since global statistics could be misleading, e.g. due to the temporal variation of the global population [12]. (3) Although there are studies on how microscopic node behavior would reproduce the observed macroscopic statistical properties [14], knowledge of how individual behaviors contribute to the evolutionary property of a single group is missing.

As part of our analysis in this paper, we propose a simple growth model, namely the active attraction (AA) model, that tries to go beyond PA by taking into account heterogeneity and local network properties in the online world. We show that this active attraction model (AA) captures both group-specific and macroscopic observations over all size ranges — as opposed to just the tail for large groups [1] or groups' initial growth and development. Our findings suggest that heterogeneity and locality play a crucial role in the dynamics of online extremist support. We also derive approximate expressions for the group size distributions in two simple systems that involve simultaneously the mechanisms of group joining (governed by either PA or AA), group leaving, and account banning, and show how these processes influence the group size distributions.

The outline of the paper is as follows. In Section 2 we introduce the dataset and the findings. In Section 3 we model the system by a PA mechanism, and check if the group-specific and macroscopic observations can be reproduced by the model. In Section 4 we introduce the AA model, and show how the group-specific and macroscopic observations are reproduced by the model. In Section 5, we derive analytic expressions for the group size distributions of two simple systems involving simultaneously the group joining (governed by either AA or PA), leaving, and banning processes. A discussion of how this work can be extended, its current limitations, and the main conclusions so far, is given in Section 6.

2. Empirical data and analysis

Our dataset is assembled using the same methodology as Ref. [1]. Our methodology for identifying these pro-ISIS online groups was as follows. We started by manually identifying relevant narratives using hashtags in multiple languages, e.g. #isn #khilafah #fisyria, and traced these to underlying online groups. The specific criterion for inclusion in the list was that the group explicitly expressed its support for ISIS, publishing ISIS-related news or propaganda and/or calling for jihad in the name of ISIS. This list was fed into software Application Programming Interfaces (APIs) that expanded it by means of automated search snowballing. The expanded online group list was then cross-checked manually to eliminate false identifications. New embedded links were manually searched to identify more online groups and hashtags. We then iterated this process until closure of the online group list (i.e. the search led back to online groups that were already in the list). Although labor intensive, we were able to find closure on a daily basis in real time. The manual list of online groups for each day was assembled to include only those appearing to express a strong allegiance to ISIS. We updated this manual list every day by logging into VK.com on a daily basis at the same time of the day, and manually searching for newly created online groups by a) analyzing posts and reposts within the known online groups; b) following selected profiles that actively publish ISIS news as well as analyzing the online groups they followed, if any. Whenever a new online group name was found, it was analyzed to establish the relevance of its narrative content and then – if appearing to be relevant as judged by the subject-matter experts to the best of their abilities – that online group was included in the database. Once the online groups supporting ISIS were identified in this way, an additional search using their followers and online groups to whom they linked was performed on that same day. Thus, the manual content analysis helped identify new online groups created – and also those that had been shutdown. An example of this manual list of online groups is available on request from the authors. Though this list is subject to human error as with any human task, the initial manual approach in our methodology using subject matter experts was crucial since online groups are a commonly occurring phenomenon on sites like VKontakte, but the vast majority of these online groups have nothing to do with pro-ISIS activity and so the pro-ISIS ones can be lost in the sea of otherwise innocuous online groups. For example, when an account page is created on VKontakte etc., there is an option to choose a ‘group’ as the category. Often these are used for informal organizations or communities of people discussing a particular sport event or concert. Unless the pro-ISIS online groups are noticed, and the finding agent (predator) shuts them down, they will continue to live within this sea of other online groups. Facebook almost immediately shuts down online groups if their content includes keywords linked to illegal activities, including terrorism.

Online groups are technically distinct from the conventional network science definition of hubs where a particular node (e.g. person) has a large number of links. For an online group, by contrast, a number of followers actually aggregated together without any one of them necessarily having many more links than the others. Joining (i.e. becoming a member of) an online group is very different from just casually visiting their website, since it gives the individual direct access to new contacts as well as continually updated information, know-how and propaganda, and also pulls him/her into ongoing discussions. Though we do not have precise numbers given the unknown accuracy of such declarations, we note that the majority of users who publicly declare a birthdate, place themselves in the age-range 18–30. Though we cannot prove that these are their true ages, the other photos, imagery and memes that they post do indeed appear largely consistent with that age-group. In terms of any relationship between banning on the one hand and creating a new group on the other, we have found that since most individuals are members of more than one group, if one of their groups is banned they simply recourse to another. Often, groups will have a similarly named but empty group waiting in case they are banned. We aim to study the use of such reserve groups in the future, though they are difficult to detect since they simply sit dormant until needed. The banning process is such that a group whose content may be visible on a given day T, appears on the next day (same URL) with a placeholder from the moderators saying, for example, that the group has been banned for violating the codes of VKontakte by promoting calls to violence. This makes it relatively easy to detect and timestamp the banning process in any given group that is being tracked.

In contrast to Ref. [1] that focuses on the late stage when groups have become very large, we focus here on the early growing stage when the group sizes are relatively small. The dataset used in this work comprises the 95 groups identified as being pro-ISIS [1] whose dates of first appearance are within our observational period of 320 days. These provide us with detailed information about the evolution of the group memberships with a high temporal resolution of a single day. The main reason why we chose this observational period is that most of the pro-ISIS groups were alive and active during that period, but became banned toward the end of the year by the moderators. After being banned, the group information becomes inaccessible. Even the groups that remained alive, tended to become inactive. This is most likely a manifestation of the migration of pro-ISIS supporters to Telegram as they became pursued with increasing vigor on VKontakte during 2015. Telegram, because of its encryption, is a safer environment for them. Our observations of the online activity and content analysis suggests that this is sufficient, because this period covers the entire lifetime of most of the pro-ISIS groups observed. The majority of groups that persist past this date (i.e. 12 out of 19) are either inactive or have dramatically changed their narrative and imagery, e.g. to include pictures of kittens playing and landscapes as opposed to the usual pro-ISIS material. Including them does not change our conclusions. Most importantly, no new pro-ISIS groups are being created. In short, the entire pro-ISIS organism was close to death on VKontakte after 320 days. Instead, our present paper focuses on the early state of group growth.

Specifically, there are three main processes involved in the group evolution: the group joining and leaving events that may occur every day during a group’s lifetime, and the banning of a group by the moderators. For each group, the dataset

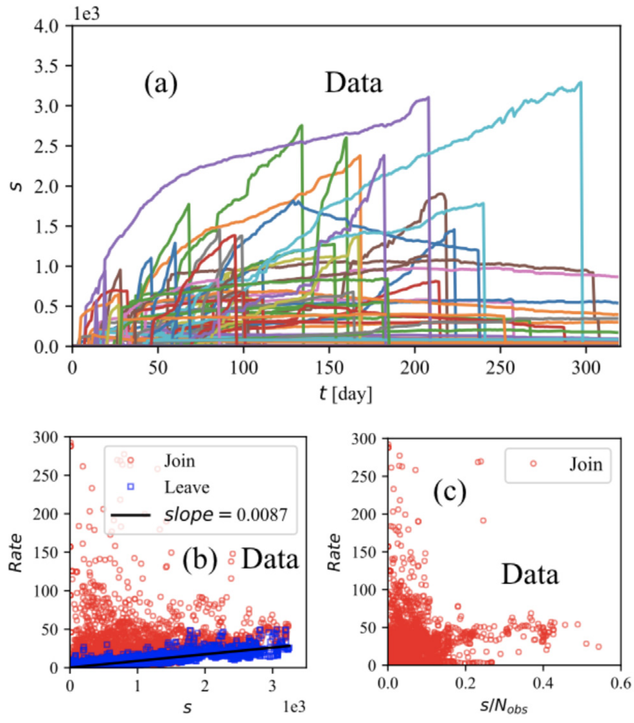

Fig. 1. (a) Time-dependent evolution of the size s of all the online groups in our empirical dataset that expressed extremist support. The vertical scale shows the size s in units of thousands (10^3). The abrupt drop of a group's size to zero indicates the group being banned by the moderators at that instant. (b) Group joining and leaving rates vs. the group sizes from the empirical data. For the leaving rates vs. the group sizes, we also show the linear regression result. (c) Group joining rates vs. group sizes normalized by the empirical global total number of follows, N_{obs}.

provides information about the size, the number of joining and leaving events on each day, as well as the first appearance date which we take as being the first day on which the group has at least one member, and the banning date if the group gets banned within the period of observation. This banning is accompanied by an abrupt drop in the group size to zero. It is possible that users could in principle be a member of all the groups observed. However, after checking our dataset, we found only a very small portion of users (0.01%) are members of all the groups during the observational period. We also note that since the groups can vary quite considerably in their focus and also their language (e.g. some are focused on pro-ISIS support in Europe, others are less geographic and focused on financing or weaponry) it makes sense that any given group is likely not of equal interest to all users,

In approaching the analysis of this data, we started by examining the evolution of the group sizes. According to PA, the size of an early-growing group should be small. However, we found that the group sizes can become very large within a few tens of days, causing a growth shape that resembles a shark-fin because of its distinctive concave shape (Fig. 1(a)). We then studied the correlation between the group joining and leaving rates on day t and the group size on that day (Fig. 1(b)). For a given group, the joining/leaving rate on day t is estimated by applying a linear fitting to the cumulative number of joining/leaving events during the 5 days around day t. We carried this out for all groups and all days — except for the first and last 2 days, since we need 5 days around day t to do the fitting. We checked that our findings are insensitive to the precise value of this time window. We found that the group-joining data points are highly dispersed when the group size is small (we consider a group is small if its size is less than $\sim 10^3$ following Ref. [1]). This suggests that non-PA rules may apply during the early growing stage. Group-leaving data points can be well fit by a line whose slope is found to be $\sim 0.0087 \pm 0.0004$ from linear regression including the standard error (Fig. 1(b)), indicating that leaving a group is more like an independent personal decision than the act of joining a group.

However, this is still insufficient to exclude PA as the governing mechanism for the group joining process during the early growth stage, for the following two reasons: First, a rapid growth of the (global) total number of group members or followers (N_{obs}) may still result in decelerating growth during the early growth stage [12]. Second, it could be that the temporal fluctuations in the total number of followers caused the dispersion of data points observed in Fig. 1(b). In order to exclude the influence of N_{obs}, we also plot the group joining rates versus the group size normalized by N_{obs} (Fig. 1(c)). It can be seen that a similarly large dispersion is also observed. To further confirm that PA is insufficient to reproduce not only the group size evolutions but also other global statistical results, in the next section we do simulations by assuming that PA governs the group-joining process.

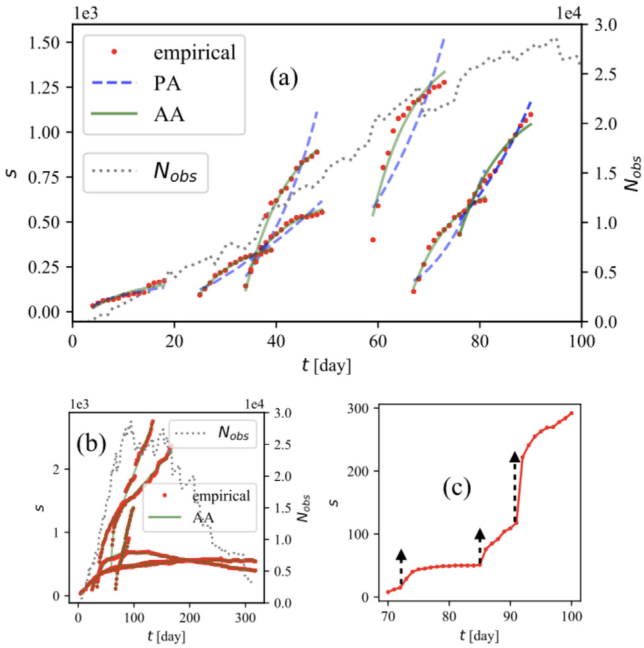

Fig. 2. (a) Model fittings for the size evolution of 7 representative extremist online groups. PA and AA correspond to fitting using Eqs. (1) and (4) respectively. Legends in (a) are applicable to all other sub-figures. For each group in (a) only the first 20 days on which the group size is nonzero are fit. (b) shows how well the groups in (a) are fit by the AA model for their entire lifetime. The abrupt drops of the group sizes to zero in (b) are due to the group being banned by the moderators. In (a)–(b), N_{obs} is the evolution of the empirical global total number of followers. (c) An example of the stair-like pattern of a group's size evolution profile. The instants when the group is exposed to a new population of potential members are marked by the dashed black arrows.

3. Preferential attachment (PA) model

3.1. Size evolution of a single group

We first model the size evolution of a single group. According to PA, the increase of the group size on day t (before the group gets banned) is given by

$$s[t] - s[t-1] = \frac{\alpha s[t-1]}{N_{obs}[t-1]} - \eta s[t-1], \tag{1}$$

where α and η describe the group joining and leaving rates, respectively, and $N_{obs}[t]$ is the total number of followers/members in the system on day t (shown in Fig. 2(a)–(b)). We define t_0 as the group's appearance date, i.e., the first day that the group's size is nonzero. Given α and initial size $s[t_0]$, η is estimated to be approximately 0.0087 (Fig. 1(b)). Since $N_{obs}[t]$ can be directly obtained from the data, we can iteratively estimate its size on all future days. Hence the curve fitting problem is to find the optimal α and $s[t_0]$ that minimize the Pearson's χ^2:

$$\underset{\alpha,\ s[t_0]}{\arg\min} \sum_{t=t_0}^{t_b} \left(\frac{s[t] - s'[t]}{s'[t]} \right)^2, \tag{2}$$

where t_b is the day when the group gets banned, $s[t]$ is the group size on day t estimated using the iterative expression (i.e. Eq. (1)), and $s'[t]$ is its corresponding observed value. For clarity, we stress that $s'[t]$ refers to an observed group size on day t while $N_{obs}[t]$ is the total number of followers/members in the system on day t. The minimization can be carried out by conventional multi-variable optimization algorithms. Note that Eq. (1) is only valid during the early growth stage when the saturation effect (i.e. the effect of finite population) can be ignored. Hence to do the fitting for each group in the dataset, we only used the sizes of the first 20 days when the group size is nonzero. The reason why we use 20 days for the growing profile of a single group is that we are interested in the early-stage growth mechanisms. The longer-term evolution of a group may involve many other factors not considered by our model that influence the evolution, such as the irregular banning activity by the moderators, or possible escalations in threats of banning, or some one-off major external event in the news, or a series of group events that happen to occur in some irregular manner. Our results show that our model fits the growth profiles of the groups well up to around the first 15–30 days after the first appearance of the groups

We find that the best estimates of α are spread across a broad range, but the fittings for most of the groups are still poor (e.g. see Fig. 2(a) where the fittings can be seen to barely reproduce the concave shapes of the empirical profiles).

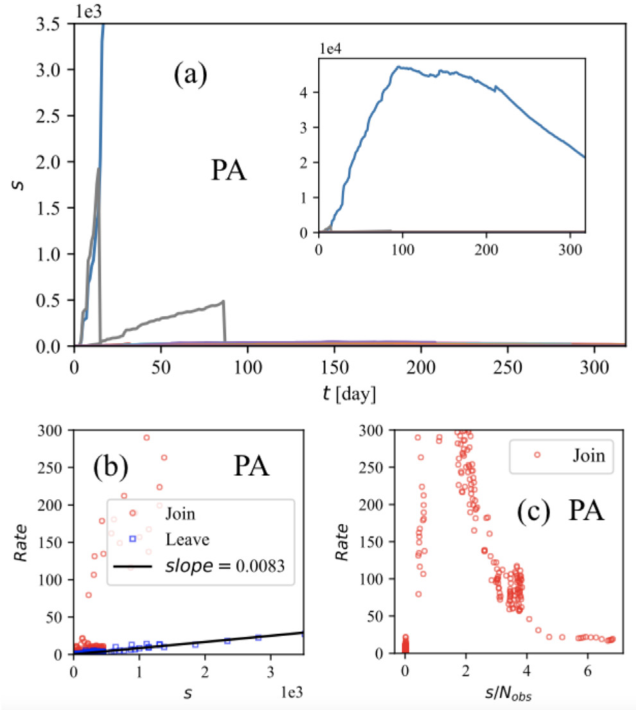

Fig. 3. (a) Time-dependent evolution of the size, s, of groups in the simulation of the PA model, where the abrupt drops of the group sizes to zero are due to the group being banned by the moderators. (b) Group joining and leaving rates vs. group sizes (s) from the simulation. For the leaving rates vs. the group sizes, we also show the linear regression result. (c) Group joining rates vs. group sizes normalized by N_{obs}.

This indicates that the empirical growth rate of $N_{obs}[t]$ is far from sufficient to bring about the observed shark-fin shapes. Also against the PA as a model mechanism is the fact that, rigorously speaking, conventional PA models implicitly assume a constant α for all groups.

3.2. Stochastic simulation of PA

We simulated the group growth using PA as follows. We set all the parameters (including the total number of groups, the creation and banning date, the total number of new joining events on each day, etc.) to be the same as the data, except that

1. we redistribute the new group joining events observed on each day to all the non-banned groups following the PA rule (i.e. the probability that a user joins a group on day t is proportional to the size of the group on day $t-1$, $s[t-1]$);
2. we use the constant group leaving rate (~ 0.0087) estimated from Fig. 1(b);
3. for a group that appears for the first time on day t_0, we manually assign a small initial size $s[t_0-1]$ to it, e.g. 1. We also tested other larger values, but the main results are left unchanged.

The exact steps in the simulation are as follows. On day t, we first detect from the dataset which groups exist (i.e. have at least one member) and denote them by a set, $\mathcal{G}$. We also obtain directly from the dataset, the total number of new joining events ($\Delta N_J[t]$). Next we redistribute the new joining events to the alive groups by drawing a sample from the multinomial distribution, *Multinomial*($\Delta N_J[t]$, $\boldsymbol{W}$), where $\boldsymbol{W}$ is a vector of probabilities whose values sum up to 1 and are proportional to the sizes of the groups on the previous day. By so doing, we ensure that the number of new joining events for a group is always proportional to its current size, and the total number of new joining events is the same as the data. Finally for each group, we consider the group leaving events by subtracting a value from its previous size. The value is drawn from the Binomial distribution *Binomial*($s[t-1]$, 0.0087), where $s[t-1]$ is the size of the group on the previous day.

Based on the simulation results, we are also able to compare the probability density function (PDF) of the group sizes in the simulation to that in the empirical observation. The PDFs for both cases are obtained as follows. First we record the sizes of all the groups on each day from the 80–100th day, to give a list S. We choose these days when doing the statistics because they correspond to a reasonably mature stage of the whole system, and hence have the maximum total number of groups. Since there are only 95 groups in our observation, by choosing such a time window, we find that we obtain better precision this way when doing the statistics. We also tested other choices of days around this we and also tried reducing the number of days, but the results are similar. Then we use the *Fit* module in the well-known *Python* package *powerlaw*[19] to obtain the PDF.

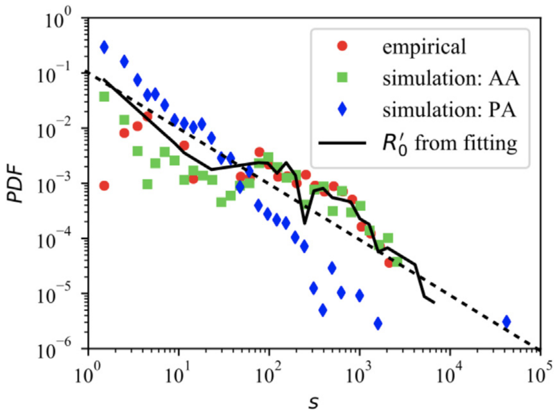

Fig. 4. Comparison between the probability density functions (PDFs) of the group sizes from the empirical observations and the simulations of the AA and PA models, and the fitting parameter R_0' estimated from the curve fittings (from Eq. (4)), for all the groups in the empirical dataset. The black dashed line has a slope of -1, and serves simply as a guide to the eye.

We find from this PA simulation that (1) the growths of most of the groups that are created at a later point in time, are effectively suppressed by a couple of groups that were created earlier (Fig. 3(a)). (2) The group-joining data points in Fig. 3(b) or (c) are much less dispersed than those in the empirical observation (Fig. 1(b) or (c)). (3) The PA-based size distribution does not agree well with the data (Fig. 4). In short, neither the group-specific nor the macroscopic observations can be well reproduced by a PA model.

4. Active attraction (AA) model

4.1. Size evolution of a single group

We now introduce the model that we refer to as AA (Active Attraction) and show that it reproduces features of the data in a way that the preferential attachment (PA) model does not. The AA starts by making two reasonable assumptions:

1. Public group events (e.g., a rapid sharing of an interesting post, and a paid advertisement or an invitation letter simultaneously sent to a large population) can attract users to join quickly even when the group size is still small [20]. Content analysis of our empirical pro-ISIS dataset, provides support for this first assumption: for example, the group with ID 276095988 rose in size from 49 to 155 on March 6 2015, just as information circulated that the ancient city of Nimrud in Assyria had been partly destroyed by ISIS. The group with ID 495490378 rose in size from 25 to 309 on February 12 2015, just as information circulated that the town of Khan al Baghdadi had been captured by ISIS. The group with ID 385900968 rose in size from 30 to 166 on February 7 2015, as information circulated that an American female hostage had been killed by ISIS and that ISIS were trying to suggest her death was due to Jordanian airstrikes. The group with ID 538079068 rose in size from 12 to 115 on January 28 2015, just after information circulated that ISIS fighters had attacked the Corinthia Hotel in Tripoli, Lybia, had killed eight people, and had then set off a car bomb.
2. Groups are heterogeneous in that a given group may not be easily accessible to all users, or it may not be of equal interest to all users. This assumption is reasonable in that a user's time is finite, and hence he/she can only consider a small fraction of all the groups that may be of interest, during his/her finite available observational period. In addition, a group can only reach a fraction of the whole population through means provided by the website such as targeted marketing [21]. Content analysis of our empirical pro-ISIS dataset, again provides support for this second assumption: for example, many of the groups are confined to particular languages (e.g. Arabic, Russian, Spanish). Also, many groups comprise members that are focused on particular geographic areas, discussing particular local news backed up exclusively by posted images from that region.

The first assumption introduces some non-PA effects, while the second stresses the heterogeneity and network locality in user-group interactions. With these two assumptions, one can imagine observing a stair-like growth in group size (e.g. Fig. 2(c)) if several public group events occur in turn. This motivates our new AA model beyond PA, and its name which reflects a more specific user-group active attraction. In particular, based on the fact that the groups in our dataset are under pressure and the topic itself will only appeal to an extreme fraction of the population of online users, a group may be accessible or of interest to only a portion of the whole population. This portion is determined by a saturation (i.e. capacity) level denoted by R. For simplicity, we assume R changes linearly with time, i.e. $R = R_0 + \gamma(t - t_0)$, where R_0 is the initial size, and γ is the rate of change. Next, we define a joining rate ($0 < \alpha < 1$ constant in time, but could be different for each group), a leaving rate (η, as mentioned earlier and found to be approximately 0.0087 for all groups), and an average (over all groups) group-banning rate β, which is only used in the simulation. With these settings, the evolution of the size of a single group before banning (hence no β dependence) in the continuous limit is given by

$$\begin{aligned}\frac{ds(t)}{dt} &= \alpha[R - s(t)] - \eta s(t) \\ &= \alpha[R_0 + \gamma(t - t_0) - s(t)] - \eta s(t),\end{aligned} \tag{3}$$

and $s(t_0) = 0$. We further define $\alpha' = \alpha + \eta$ as the effective joining rate, and $R' = \alpha R/(\alpha + \eta)$ as the effective saturation level. Let R'_0 be $\alpha R_0/(\alpha + \eta)$, and γ' be $\alpha\gamma/(\alpha + \eta)$. Then the differential equation can be rewritten as $ds(t)/dt = \alpha'[R' - s(t)] = \alpha'[R'_0 + \gamma'(t - t_0) - s(t)]$, whose solution is given by

$$s(t) = (R'_0 - \frac{\gamma'}{\alpha'})[1 - e^{-\alpha'(t-t_0)}] + \gamma'(t - t_0), \tag{4}$$

which contains 4 free parameters (i.e., t_0, R'_0, α' and γ'), and is the expression used for the fittings by the AA model shown in Fig. 2(a)–(b). The fitting is performed using a conventional multi-variable optimization algorithm. We used the *curve_fit* function with bounds set properly in the well-known *Python* package *SciPy*[22].

We find that Eq. (4) fits most of the shark-fin group growth profiles well, not only in terms of their early growth stage (e.g. Fig. 2(a)), but also over their entire lifetime in many cases (e.g. Fig. 2(b)). In addition, we find that there exists a high heterogeneity in the fitting parameters. That is, R'_0, α' and $|\gamma'|$ (γ' can be either positive or negative) all spread over a broad range — and indeed, they scale roughly like a power-law with the power-law exponent ~ 1, except t_0 which is not very sensitive and is always just a couple of days earlier than the observed first-appearing day. We also tried constraining some or all of the three parameters (i.e., R'_0, α' and γ') to a set of reasonable constant values (e.g., the median or mean values from previous fittings in which they had been allowed to change in time), but we found that this resulted in poorer fittings. Hence, for the empirical data, the heterogeneity indeed exists and manifests itself in all these three parameters. Since the distributions of these three parameters are not the focus of this work, we show only the distribution of R'_0 in Fig. 4, which is used in the simulations.

4.2. Stochastic simulation of AA

We now show how we performed the simulation of the AA model. Similar to the simulation of the PA model, we controlled all the parameters to be the same as the dataset, except that (1) we adopt a constant group-leaving rate of ~ 0.0087; (2) we redistribute the new joining events on each day among all the alive groups on that day according to the AA rule (i.e. the probability that a user joins a group on day t is proportional to $R[t-1]-s[t-1]$, where $R[t-1]$ and $s[t-1]$ are the saturation level and the group size on day $t - 1$, respectively); (3) we further assume $\gamma = 0$ for each group (and hence R is constant), and set the saturation level of each group to R_0 estimated from the curve fitting by Eq. (4). Note that the AA rule we used implicitly assumes α is a constant for all the groups – otherwise the probability of joining a group should be proportional to $\alpha(R[t - 1] - s[t - 1])$ – and hence the simulation focuses on studying the effect of the heterogeneity in the saturation level. We could have also set the α of each group to be the value obtained from the fitting, which would make our simulation agree even better with the dataset, but our simulations show that the heterogeneity in R_0 is sufficient to reproduce the empirical observations. To make our comparison between AA and PA as meaningful as possible, we performed the AA simulations in a similar way to our PA simulations, as follows: On day t, we first detect from the dataset which groups are alive (i.e. have at least one member) and denote them by a set, $\mathcal{G}$, and also obtain directly from the dataset the total number of new joining ($\Delta N_J[t]$) events. Next we redistribute the new joining events to the alive groups by drawing a sample from the multinomial distribution, *Multinomial*$(\Delta N_J[t], \boldsymbol{W})$, where $\boldsymbol{W}$ is a vector of probabilities whose values sum up to 1 and are proportional to $\alpha(R[t - 1] - s[t - 1])$. Finally for each group we consider the group leaving events by subtracting a value from its previous size, using a value drawn from the Binomial distribution *Binomial*$(s[t - 1], 0.0087)$.

Even though the heterogeneity in α and γ is neglected, we find that the AA simulation (Fig. 5) reproduces well the empirical shark-fin shapes (Fig. 1(a)). It also reproduces the dispersion patterns in the joining and leaving rates (Fig. 1(b) and (c)), as well as the group size distribution (Fig. 4).

5. Theoretical analysis of the group size distribution

We now study mathematically the group size distribution of the system involving simultaneously the three mechanisms of group joining, leaving, and banning due to moderators. Since similar systems can exist in many other areas of human extremist activity, we cast our discussion in a more general sense and show the analytical results for both the case when the group joining follows the AA rule, and separately when it follows the PA rule. Conventional rigorous treatment of the problem may involve solving the master equations, but that approach becomes intractable in practice — especially when taking into account the heterogeneity of the coefficients. Hence, we resort to an approximative method and study only two simple cases.

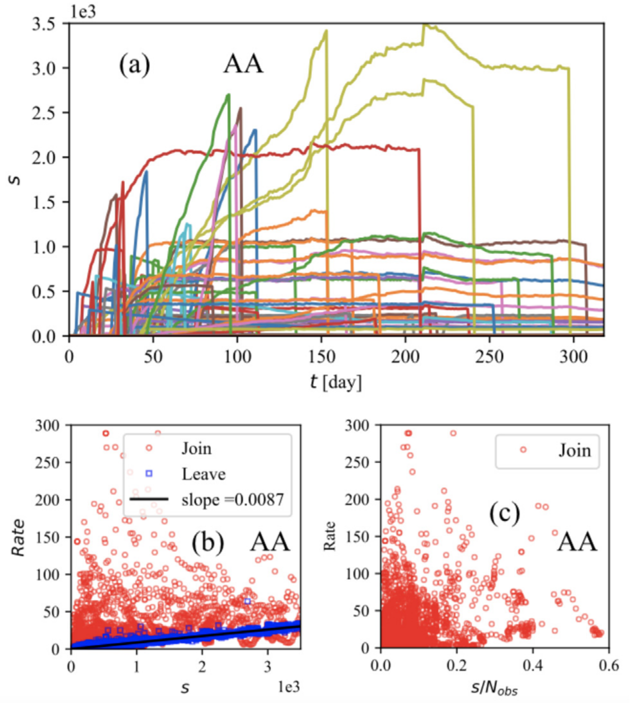

Fig. 5. (a) Time-dependent evolution of the size, s, of all the groups in the simulation of the AA model, where the abrupt drops of the group sizes to zero are due to the group being banned by the moderators. (b) Group joining and leaving rates vs. group sizes (s) for the simulation. For the leaving rates vs. the group sizes, we also show the linear regression result. (c) Group joining rates vs. group sizes normalized by N_{obs}.

5.1. The AA case

We start with the case of the group joining being an AA process. For simplicity, we consider the case that the groups are created with the same probability on each day during the periods of observation; hence the probability that a group created on day t_0 will remain alive (or observable) on day t is

$$P_{T_\Delta}(t_\Delta) \sim (1-\beta)^{t_\Delta}, \tag{5}$$

where β is the probability that a group will get banned on a given day, and $t_\Delta \equiv t - t_0$. As mentioned earlier, in the case that the joining is an AA process then the size evolution of a group, $s(t)$ (hereafter, we omit its t label for convenience), is governed by $ds/dt = \alpha'(R' - s)$. For simplicity, we consider the case when R' and α' are both time-independent. In such a case, the solution of this differential equation when the initial condition $s(t_0) = 0$ is given by $s = R'(1 - e^{-\alpha' t_\Delta})$, from which we can inversely get $t_\Delta = -\ln(1 - s/R')/\alpha'$, and hence $dt_\Delta/ds = 1/[\alpha'(R' - s)]$. Consider the distribution of R'_0 in Fig. 4, the detailed shape of which (though not the main concern of this work) is irregular – probably due to poor statistics or the irregular banning events of the moderators etc. – but which is very roughly distributed around a power-law of exponent ~ 1. Inspired by this shape, we consider the simple case that the effective saturation level R' of all the groups follows a power-law distribution of $P_{R'}(R') \sim R'^{-\lambda}$, and that a' is the same for all groups. Then by ignoring the stochastic fluctuations in the group size evolution profiles (meaning that there is now a one-to-one correspondence between t_0, or t_Δ, and s for a given R' and t), Eq. (5) becomes after a change of variables:

$$P_S(s|R') \sim (1-\beta)^{-\ln(1-s/R')/\alpha'}(dt_\Delta/ds). \tag{6}$$

Following the definition of R', which is the effective saturation level, we know that $R' \geq s$ when assuming R' is time-independent. Hence the probability density that a group is observed at t having size s (or created at t_0) is given by

$$\begin{aligned} P_{AA}(s) &\sim \int_s^\infty P_S(s|R')P_{R'}(R')dR' \\ &\sim \int_s^\infty (1-\beta)^{t_\Delta} P_{R'}(R')(dt_\Delta/ds)dR' \\ &\sim \int_s^\infty (1-s/R')^{-\ln(1-\beta)/\alpha' - 1} R'^{-(\lambda+1)} dR'. \end{aligned} \tag{7}$$

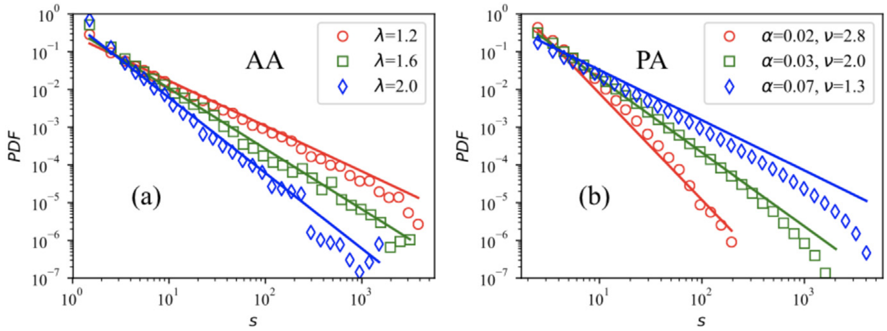

Fig. 6. Group size distributions for the (a) AA and the (b) PA model. Plots show simulation results (dots) vs. analytic results (solid lines, which have the same color as the corresponding simulation result). Results for different parameters are shown.

Expanding the integrand with respect to s/R' and keeping the terms up to $O(s/R')$, we obtain

$$P_{AA}(s) \sim \int_s^\infty \left\{1 + \left[\frac{\ln(1-\beta)}{\alpha'} + 1\right]\frac{s}{R'}\right\} R'^{-\lambda-1} dR' \\ \sim s^{-\lambda}, \tag{8}$$

i.e. it also follows a power-law distribution that scales approximately the same way with the distribution of R'. This similarity between the saturation level distribution and the group size distribution is also found from the empirical data (Fig. 4).

5.2. The PA case

For comparison, we also calculate in a similar way the group size distribution for the case that the group joining is a PA process. Neglecting the finite-population effect and considering only the case when $N(t)$ is a constant, we have $ds/dt \sim (\alpha - \eta)s$, whose solution $s \sim e^{t_\Delta/(\alpha-\eta)}$. Hence $t_\Delta \sim \ln(s)/(\alpha - \eta)$ and $dt_\Delta/ds \sim 1/[s(\alpha - \eta)]$. Ignoring the stochastic fluctuations in the group size evolution profiles, there is a one-to-one correspondence between t_Δ, or t_0, and s for a given t. Therefore, through a change of variables, the probability density of the group size is given by

$$P_{PA}(s) \sim (1-\beta)^{t_\Delta} dt_\Delta/ds \sim s^{-\nu} \tag{9}$$

where $\nu = -\ln(1-\beta)/(\alpha - \eta) + 1$, which is greater than 1 for any $0 < \beta < 1$ and $\alpha > \eta$.

5.3. Numerical verification and discussion

We now check these analytical results since their derivations involve several approximations. We conducted stochastic simulations for both systems. The simulation settings comply with the basic assumptions that we made when defining the systems for the analytic derivations. For the AA case, we initialize 1000 groups by assigning to each of them an initial group size of 1 and a saturation level (R) sampled from a Zipf's distribution with a power-law exponent equal to λ. Next, on each day t we generate the number of joining events for each group by sampling from the binomial distribution $Binomial(R-s[t-1], \alpha)$ and add it to $s[t-1]$ to obtain $s[t]$, where α is the joining rate. Then we generate the number of leaving events for each group by sampling from $Binomial(s[t-1], \eta)$, where η is the leaving rate which was earlier determined to be 0.0087. Finally we ban each group with a probability β (without loss of generality, we use a value of 0.02) by setting the number of members to 1, which means that after banning we immediately create another group so that the total number of groups is always a constant.

We run the simulations for a sufficiently large number (e.g. 2000) of steps (i.e. days) for each run. This means that the creation dates of the groups are effectively randomized automatically after a long period due to the randomness in the group banning, and the total number of followers will also converge to a stable level. In addition, to improve the precision, we repeat the simulation 10 times for every set of parameters and then obtain the average distributions by combining the sample from each run. For each run, without loss of generality, the sample is formed by the group sizes on each day during the 80–100th day of the last 320 days in the simulation. With respect to the simulation for the PA case, everything is the same as for the AA case, except that we do not need R' anymore. On each day t, the first step in the AA case (see above) is to generate the number of joining events for each group. The number is sampled from the binomial distribution $Binomial(s[t-1], \alpha)$ and is added to $s[t-1]$ to obtain $s[t]$.

Fig. 6 shows that our analytical results agree well with the stochastic simulations. We can see that for the AA model, the distribution of the saturation level plays a pivotal role in determining the group size distribution. By contrast, for the PA model all three mechanisms matter. In addition for the PA case, when the group-banning rate is small and the group-leaving rate is significantly smaller than the joining rate, the power-law exponent is always around 1. We can also see that different microscopic user-group interaction mechanisms (e.g. AA and PA) may result in the same power-law exponent in the macroscopic statistics, which means that how a macroscopic quantity scales can be insufficient to infer the microscopic mechanism. Hence user and group level analysis – as presented in this paper – become important.

6. Discussion and conclusions

We have shown that the non-PA effect is crucial for explaining the rapid growth of groups observed during the early growth stage in a system of online extremist support. The PA effect then becomes more appropriate when the group size is large, as indicated in Ref. [1]. We proposed a simple non-PA model, which we refer to as the AA model, that catches the effects of locality and microscopic heterogeneity in the dynamics of group formation and which are overlooked by conventional PA models. While we concede that there could be alternative models, our AA model is arguably one of the simplest, and it explains well both the group-specific and global statistical observations. Since the group joining probability in the AA process is not positively correlated to the current size of a group, the AA model provides a possible explanation for the cascading joining that results in a shark-fin shape observed in individual groups' size evolution. Such an observation can barely be described by a PA model without introducing an unusually high growth rate for the global total number of followers. In addition, since both PA and AA could produce a similar group size distribution, this work suggests the importance of a deeper understanding of the behaviors of individual users and groups.

Building on the work in this paper, one possibility for future research would be to quantify the extent to which the dynamics for pro-ISIS fighting groups in the real world, correlate with the dynamics reported here of the pro-ISIS online groups. Though there is insufficient data available at the granular level to carry out a thorough scientific investigation now, such group-level conflict data may indeed become available in the future. However, there are a number of things that we can already say, based on our preliminary investigations, that support the notion that there is a relationship between pro-ISIS fighting groups in the real world, and pro-ISIS groups in the online world. First, there are a wealth of individual examples within the content of the groups, that suggest that the groups themselves are involved in preparations of certain real-world operations: for example, there was a March 8, 2015 posting within several groups saying '[translation] ISIS are preparing to attack the city of Karbala, 500 tons of explosives are ready' which was then followed several months later by the discovery of booby-trapped vehicles and ISIS members in a small town 80 km west of Karbala. Second, we also found evidence of real-world activity being discussed in real time in these online pro-ISIS groups. Here are some examples, with specific web-links available from the authors for each: (1) Evidence of fundraising: Multiple incidents of collecting funds for potential fighters who want to travel to Syria but cannot afford it. Also transfer of funds for fighters who are already in Syria. (2) Evidence of real-time operational information stream: Some online groups resemble an alternative news outlet where they stream information directly from their territory. These include operational updates from the battlefield, e.g. the specifics of the Kobane radio tower in real-time. One example image says "..ISIS took control over the Kobane's radio tower ... Mujahedeen advanced 500 meters into Kobane". (3) Evidence of mobilizing support: Images include text such as: "Brothers! Yesterday, in a German town of Celle, a 100-people mob of Yazidi Kurds beat up 5 Chechens in retaliation for Chechens fighting within ISIS in Iraq and Syria where they kill Kurds. Since Celle has the largest Yezidi Kurds community (about 5000 people), our local brothers' lives are under threat. This is a call to all brothers from nearby locations to send groups of 30–40 people to protect our brothers in distress!" Actual news reports of the event show that the online group had not invented the story, and hence the online group was indeed operationally relevant. Another screenshot says: "Urgent! Those who are traveling to Celle for the fight, do NOT go! The Kurds have apologized and there will be a peaceful settlement. German border authorities are not allowing Chechens enter from Austria, France and other countries. All German cities are full of police patrols". It was published a day after the previous one. (4) Teaching survival skills: Some pro-ISIS online groups include advice on cellphone and Internet use during an operation in order to avoid being detected by security services; and also ways to prevent or repel a drone attack during an operation. (5) Recruitment: Evidence that the online groups serve as a platform to spread recruitment messages, is illustrated by an example available from the authors that starts with the statement: "IS fighters in Dagestan call other Caucasus mujahedeen to enter their ranks". Indeed many Caucasus guerrilla groups later joined ISIS. We have also found suggestion of a connection using data on pro-ISIS attacks from the well-known and highly respected Global Terrorism Database (GTD) which is an open-source database including information on terrorist events around the world – even though it is of course far from perfect. Specifically, the distribution of severities of pro-ISIS attacks for small values of the severity appears to crudely mimic a power-law with a slope near 1. If we suppose that events with low severities involve smaller groups of fighters rather than larger army-like sizes, and hence we suppose that the distribution of severities for smaller severity values is a proxy for the distribution of groups at smaller group size, the results are at least consistent with the findings in Fig. 4. While these results are very tentative given the nature of the empirical data, they are sufficiently encouraging to warrant more careful collection of data on ISIS fighter activity for comparison to our findings concerning online pro-ISIS groups.

There is also some evidence of a mapping between online users in these online pro-ISIS groups and real-world fighters. Specifically, we have found that particular users are sometimes referred to by members of a group in connection with actual events in the real world. By means of this connection, which typically comes through other users of a group praising a certain user (which real name and identity we would not otherwise have known), we have been made aware that members of the groups in our dataset include a mother from Sweden who migrated with her three children to join ISIS; a combatant who produced real-time audio recordings with street-level detail during assaults in Syria; a suicide bomber who then seems to have driven a truck of explosives into a Shia army in Iraq; females with 'selfie' fighter photos; someone whose combat activity was so specialized that their later death fighting in Iraq made international media headlines; and someone who transitioned from a lecturer to later become the known leader of Chechen fighters in ISIS. In terms of typical users, we chose one randomly who we refer to here as XXX Shishani: he joined about 40 online groups and produced a sequence of postings and activity show that he progressed within a matter of months from an initial broad pro-ISIS rhetoric which included an image of a killed fighter (January 11, 2015) to a focus on weaponry (January 26, 2015) and then operational details such as building smoke-screens and the mechanics of a gun (March 4, 2015) – with the net result being an individual whose profile photo presents an extremist pose in clothing consistent with a country in U.S. or northern Europe, as shown in the final photo. The first part of his posted name has Anglo-Saxon and German roots, while the other part 'Shishani' is a common family name among the Chechen diaspora in the Arab world and the United States. He only has 56 followers online, which is consistent with the idea that this person was not a special fighter or leader in any sense, but rather an inspired individual from 'out of the blue'. On 10 November 2015, a public posting for him appeared by another individual saying "[translation] God willing brother, see you in Akhirat" where Akhirat is an Islamic term referring to the afterlife. This, together with his recent lack of online activity, suggest that he may already have been killed.

Of course, there are still many open questions. For instance, the origin of the user and group level heterogeneity is still unknown. More specifically, it is not clear if the broad distribution of the saturation levels originates from the topology of the network that influences the accessibility of the groups (i.e. the locality), or is due to some other more complicated mechanisms that result in the heterogeneity in users' interests in the groups. Hopefully the present work draws researchers' attention to these open questions, and can serve as a stepping-stone toward answering them. Eventually we hope to provide a clear practical payoff from these findings by working with subject-matter experts in the policy field to evaluate actionable policies. We speculate that our findings suggest a policy approach in which the dynamics of present and future extremist online groups can be estimated and hence impacted based on a subject-matter expert's best guess as to the heterogeneity of group members (e.g. the extent to which group members use different memes or imagery) and network locality in terms of access to groups (e.g. groups using the same language). With this estimate in hand, inferences can in principle be drawn as to plausible interventions to prevent or delay their build-up. We stress that our ambitions are more modest in this paper, in that we focus here on the initial stage of establishing potentially useful details of the mechanistic process underlying the group dynamics. Work can then hopefully follow with subject-matter experts in the policy field, to explore actionable policies. We also note that there may be an interesting, and possibly strong, correlation between group size and exogenous events that is worth further examination and which would provide a compelling need to further investigate the dynamics of these online groups. We are now in the process of obtaining data about such external events, which can then be compared to group activity and content.

Acknowledgments

We are grateful to Yulia Vorobyeva, Andrew Gabriel and Anastasia Kuz for initial help with data collection and analysis. NFJ gratefully acknowledges funding under National Science Foundation (NSF), USA grant CNS 1522693 and Air Force (AFOSR), USA grant FA9550-16-1-0247. The views and conclusions contained herein are solely those of the authors and do not represent official policies or endorsements by any of the entities named in this paper.

References

[1] N. Johnson, M. Zheng, Y. Vorobyeva, A. Gabriel, H. Qi, N. Velasquez, P. Manrique, D. Johnson, E. Restrepo, C. Song, et al., Science 352 (6292) (2016) 1459–1463.
[2] L. Backstrom, D. Huttenlocher, J. Kleinberg, X. Lan, Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, 2006, pp. 44–54.
[3] A. Mislove, M. Marcon, K.P. Gummadi, P. Druschel, B. Bhattacharjee, Proceedings of the 7th ACM SIGCOMM conference on Internet measurement, ACM, 2007, pp. 29–42.
[4] G. Palla, A.-L. Barabasi, T. Vicsek, Nature 446 (7136) (2007) 664–668.
[5] J. Leskovec, K.J. Lang, A. Dasgupta, M.W. Mahoney, Internet Math. 6 (1) (2009) 29–123.
[6] J. Leskovec, K.J. Lang, A. Dasgupta, M.W. Mahoney, Proceedings of the 17th international conference on World Wide Web, ACM, 2008, pp. 695–704.
[7] A.L. Traud, P.J. Mucha, M.A. Porter, Physica A 391 (16) (2012) 4165–4180.
[8] L. Backstrom, R. Kumar, C. Marlow, J. Novak, A. Tomkins, Proceedings of the 2008 International Conference on Web Search and Data Mining, ACM, 2008, pp. 117–128.
[9] J. Berger, H. Perez, Available at https://goo.gl/wyq1C8, 2016.
[10] S.R. Kairam, D.J. Wang, J. Leskovec, Proceedings of the fifth ACM international conference on Web search and data mining, ACM, 2012, pp. 673–682.
[11] R. Albert, A.L. Barabási, Rev. Modern Phys. 74 (1) (2002) 47.
[12] A.-L. Barabási, H. Jeong, Z. Néda, E. Ravasz, A. Schubert, T. Vicsek, Physica A 311 (3) (2002) 590–614.
[13] R. Kumar, J. Novak, A. Tomkins, Link mining: models, algorithms, and applications, Springer, 2010, pp. 337–357.
[14] J. Leskovec, L. Backstrom, R. Kumar, A. Tomkins, Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, 2008, pp. 462–470.
[15] J.L. Kelley, L.J. Morrell, C. Inskip, J. Krause, D.P. Croft, PLoS One 6 (8) (2011) e24280.
[16] J. Yang, J. Leskovec, Knowl. Inf. Syst. 42 (1) (2015) 181–213.
[17] C. Sorensen, H. Zhang, T. Taylor, Phys. Rev. Lett. 59 (3) (1987) 363.
[18] S. Gueron, S.A. Levin, Math. Biosci. 128 (1) (1995) 243–264.
[19] J. Alstott, E. Bullmore, D. Plenz, PLoS One 9 (1) (2014) e85777.
[20] K. Kanin, https://goo.gl/FuYbS9, KonstantinKanin.com, DaCompany Ltd., 2017.
[21] M. Martin, https://goo.gl/xAWxz9, Cleverism, 2014.
[22] E. Jones, T. Oliphant, P. Peterson, et al., SciPy: Open source scientific tools for Python, http://www.scipy.org/, 2001 onwards, [Online: accessed on 2017-09-25].