

The Brookings Project on U.S. Relations with the Islamic World
ANALYSIS PAPER | No. 20, March 2015

The ISIS Twitter Census

Defining and describing the population
of ISIS supporters on Twitter

BY J.M. BERGER AND
JONATHON MORGAN



Table of Contents

1	The Authors	39	3. Methodology
1	The Research Team	40	3.1 Starting Point
2	Executive Summary	41	3.2 Challenges and Caveats
5	Introduction	42	3.3 Bot and Spam Detection
6	1. Demographics of ISIS Supporters	43	3.4 Description of Data Collected
	1.1 Estimating the Total Number of Supporters	44	3.5 Data Codebook
8	1.2 The Demographics Dataset	45	3.6 Sorting Metrics
9	1.3 Data Snapshot	47	3.7 Metrics Performance and Estimates
10	1.4 Location	50	3.8 Machine Learning Approach to Level 2 Results
12	1.4.1 Inferred Location	51	3.8.1 Predicting Supporters
14	1.5 Languages	52	4. Conclusions: Pros and Cons of Suspending ISIS Supporters on Twitter
15	1.6 Display Names		4.1 Intelligence Value
16	1.7 Date of Account Creation	54	4.2 Effectiveness of Suspensions in Limiting ISIS's Influence
18	1.8 Date of Last Recorded Tweet	55	4.3 Suspensions and Trade-offs
19	1.9 Avatars	58	4.4 Preliminary Policy Recommendations
20	1.10 Top Hashtags	59	4.5 Research recommendations
21	1.11 Top Links	61	62 Appendices
23	1.12 "Official" Accounts	62	A) Notes on In-Network Calculations
24	1.13 Bots and Apps	63	B) Control Group
26	1.14 Smartphones	64	About the Project on U.S. Relations with the Islamic World
27	2. Social Media Metrics	65	The Center for Middle East Policy
28	2.1 Tweeting patterns		
30	2.2 Number of Followers		
32	2.3 Number of Accounts Followed		
32	2.4 Follower/Following Ratio		
33	2.5 Suspensions		
33	2.5.1 Number of observed suspensions		
34	2.5.2 Suspensions of accounts created in September and October 2014		
34	2.5.3 Performance of accounts that were not suspended		
36	2.5.4 Partial comparison data		

The Authors

J.M. Berger

J.M. Berger is a nonresident fellow with the Project on U.S. Relations with the Islamic World at Brookings and the author of *Jihad Joe: Americans Who Go to War in the Name of Islam* (Potomac Books, 2011) and *ISIS: The State of Terror* (Ecco, 2015). An analyst and consultant studying extremism, he is also involved with developing analytical techniques to study political and extremist uses of social media. He is a regular contributor to *Foreign Policy* and the founder of Intelwire.com.

Jonathon Morgan

Jonathon is a technologist, data scientist, and startup veteran. He runs technology and product development at CrisisNET, Ushahidi's streaming crisis data platform, and consults on machine learning and network analysis. Morgan is also co-host of Partially Derivative, a popular data science podcast.

Prior to Ushahidi, Morgan served as the CTO of SA Trails, a venture-backed startup with operations across South America. He was Principal Technologist at Bright & Shiny, where he led the team that realized scientist David Gelertner's vision of personalized data streams ("lifestreams"). Previously, Morgan served as CEO of StudentPositive, an education technology company focused on predicting student behavior.

The Research Team

J.M. Berger and Jonathon Morgan designed and implemented collection and analysis techniques using proprietary code, supplemented by a limited number of third-party tools. Berger, Youssef ben Ismail, and Heather Perez coded accounts. Perez also contributed essential data points pertaining to ISIS's social network strategy. J.M. Berger wrote the report.

Special thanks are due to Will McCants, Anne Peckham, and Kristine Anderson of the Brookings Institution; Yasmin Green, Jana Levene, and Justin Kosslyn of Google Ideas; and Jessica Stern, Chris Albon, Dan Sturtevant, and Bill Strathearn.

This paper was commissioned by Google Ideas and published by the Brookings Institution. The views expressed herein represent those of the authors alone.

Executive Summary

The Islamic State, known as ISIS or ISIL, has exploited social media, most notoriously Twitter, to send its propaganda and messaging out to the world and to draw in people vulnerable to radicalization.

By virtue of its large number of supporters and highly organized tactics, ISIS has been able to exert an outsized impact on how the world perceives it, by disseminating images of graphic violence (including the beheading of Western journalists and aid workers and more recently, the immolation of a Jordanian air force pilot), while using social media to attract new recruits and inspire lone actor attacks.

Although much ink has been spilled on the topic of ISIS activity on Twitter, very basic questions remain unanswered, including such fundamental issues as how many Twitter users support ISIS, who they are, and how many of those supporters take part in its highly organized online activities.

Previous efforts to answer these questions have relied on very small segments of the overall ISIS social network. Because of the small, cellular nature of that network, the examination of particular subsets such as foreign fighters in relatively small numbers, may create misleading conclusions.

The information vacuum extends to—and is particularly acute within—the sometimes heated discussion of how the West should respond to this online campaign.

While there are legitimate debates about the bounds of free speech and the complex relationship between private companies and the public interest, some have argued against suspending terrorist social media accounts on the basis that suspensions are not effective at impeding extremist activity online. These arguments that are usually predicated

on very small samples of potentially misleading data, when data is proffered at all.

We set out to answer some of these important questions using innovative techniques to create a large, representative sample of accounts that can be clearly defined as ISIS supporters, and to attempt to define the boundaries of ISIS's online social network.

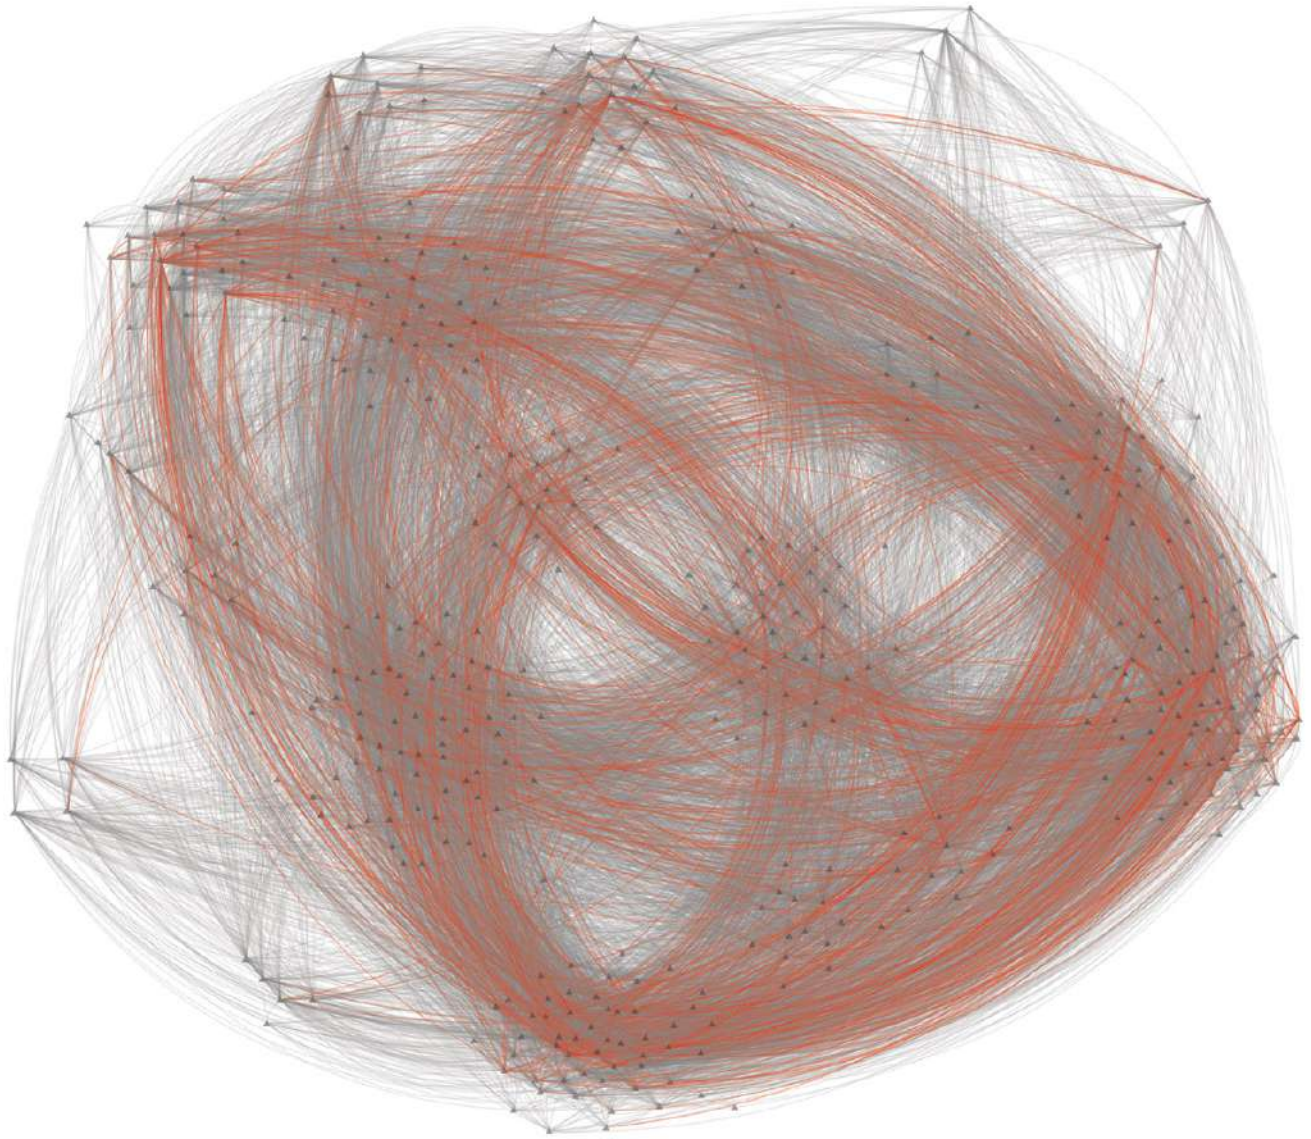
The goals of the project included:

- Create a demographic snapshot of ISIS supporters on Twitter using a very large and accurate sample of accounts (addressed in sections 1 and 2 of this paper).
- Outline a methodology for discovering and defining relevant accounts, to serve as a basis for future research using consistent comparison data (section 3).
- Create preliminary data and a path to further investigate ISIS-supporting accounts suspended by Twitter and the effects of suspensions (section 2.5).

Our findings, based on a sample of 20,000 ISIS supporter accounts, include:

- From September through December 2014, we estimate that at least 46,000 Twitter accounts were used by ISIS supporters, although not all of them were active at the same time.
- The 46,000 figure is our most conservative estimate for this time frame. Our maximum estimate is in the neighborhood of 70,000 accounts; however, we believe the truth is closer to the low end of the range (sections 1.1, 3.5, 3.6, 3.8).
- Typical ISIS supporters were located within the organization's territories in Syria and Iraq, as well as in regions contested by ISIS. Hundreds of ISIS-supporting accounts sent tweets with location metadata embedded (section 1.4).

- Almost one in five ISIS supporters selected English as their primary language when using Twitter. Three quarters selected Arabic (section 1.5).
- ISIS-supporting accounts had an average of about 1,000 followers each, considerably higher than an ordinary Twitter user. ISIS-supporting accounts were also considerably more active than non-supporting users (section 2).
- Much of ISIS's social media success can be attributed to a relatively small group of hyperactive users, numbering between 500 and 2,000 accounts, which tweet in concentrated bursts of high volume (section 2.1).
- A minimum of 1,000 ISIS-supporting accounts were suspended between September and December 2014, and we saw evidence of potentially thousands more. Accounts that tweeted most often and had the most followers were most likely to be suspended (section 2.5.1).
- At the time our data collection launched in September 2014, Twitter began to suspend large numbers of ISIS-supporting accounts. While this prevented us from creating a pre-suspension dataset, we were able to gather information on how the removal of accounts affected the overall network (section 2.5.4).
- Account suspensions do have concrete effects in limiting the reach and scope of ISIS activities on social media. They do not, at the current level of implementation, eliminate those activities, and cannot be expected to do this. Some critics argue suspensions are ineffective because ISIS propaganda is still available on Twitter. Any balanced evaluation of current levels of suspension activity clearly demonstrates that total interdiction is not the goal. The qualitative debate is over how suspensions affect the performance of the network and whether a different level of pressure might produce a different result (sections 2.5, 4.2). While it is possible to target suspensions in a manner that would be far more devastating to ISIS networks, we do not advise such an approach for several reasons (sections 4.1 and 4.3).
- The process of suspension does create certain new risks. Most importantly, while suspensions appear to have created obstacles to supporters joining ISIS's social network, they also isolate ISIS supporters online. This could increase the speed and intensity of radicalization for those who do manage to enter the network, and hinder organic social pressures that could lead to deradicalization (section 4.3).
- Further study is required to evaluate the unintended consequences of suspension campaigns and their attendant trade-offs. Fundamentally, tampering with social networks is a form of social engineering, and acknowledging this fact raises many new, difficult questions (section 4.3).
- Social media companies and the U.S government must work together to devise appropriate responses to extremism on social media. Although discussions of this issue often frame government intervention as an infringement on free speech, in reality, social media companies currently regulate speech on their platforms without oversight or disclosures of how suspensions are applied (section 4.4). Approaches to the problem of extremist use of social media are most likely to succeed when they are mainstreamed into wider dialogues among the wide range of community, private, and public stakeholders.



Links among the top 500 Twitter accounts as sorted by the in-group metric used to identify ISIS supporters. Red lines indicate reciprocal relationships.

Introduction

This study consists of four parts:

- ISIS supporter demographics
- ISIS supporter social media metrics
- A detailed discussion of the methodology used for this paper
- A preliminary examination of the effects of suspending social media accounts and recommendations for further study and policies

The first two sections are based on a sample of 20,000 accounts believed to be comprised of at least 93 percent ISIS supporters. We examine where these supporters are located, what languages they speak, what identifying information they provide, when their accounts were created, a limited view on what content they post, and the methods they use to spread ISIS propaganda and recruit followers around the globe.

The third section discusses in considerable detail how we identified these accounts. We believe this is a crucial part of the discussion, to allow readers to determine how much confidence to place in the results, and to establish a framework for future research on the performance of social networks.

The fourth section discusses some of the implications and questions raised by this study, particularly pertaining to the effects of suspending extremist social media accounts. This section also points to some of the challenges of designing a coherent approach among all stakeholders involved in countering the problem of violent extremism on social media.

1. Demographics of ISIS Supporters

Who are the users supporting ISIS on Twitter? Where do they live, and how do they do their work online? These are questions of great importance for anyone trying to understand the scope of the problem and possible remedies.

Using a variety of innovative approaches, we identified an accurate dataset of 20,000 ISIS supporter accounts on Twitter. Through this sample we aimed to estimate the total number of accounts supporting ISIS on Twitter and to create a demographic profile of this group, shedding light on where users are based, what languages they speak, what they tweet about, and how they access the Internet.

The demographics data, as well as the Twitter metrics discussed in section 2, also shed light on the social media strategies ISIS uses to disseminate its messages online.

1.1 Estimating the Total Number of Supporters

We had hoped to establish both a floor and a ceiling for estimates of the size of ISIS's supporter base on Twitter. A completely reliable ceiling proved elusive due to the size of the dataset, its rapid evolution, and the complexity of the relationships within it. However we were able to establish a floor with reasonable certainty.

During the period of October 4 through November 27, 2014, we estimate there were no fewer than 46,000 Twitter accounts supporting ISIS. Some accounts that were active in September but subsequently suspended were relevant to the set, and different kinds of information were collected at different speeds based on Twitter API limits. This figure excludes deceptive tactics meant to inflate ISIS's Twitter following, such as automated bots, but includes multiple accounts maintained by human users.

This estimate was derived from two sources of data:

- We collected extremely robust data on nearly 50,000 accounts. We estimate that a minimum of 30,000 of these are accurately described as accounts belonging to ISIS supporters and controlled by a human user, using the most conservative criteria. We have a high level of confidence in this estimate, which is based on samples coded under the criteria described in section 3.5 and the metrics described in sections 3.6 and 3.7.
- We also collected partial data on 1.9 million additional accounts, as described in 3.8. Because this data was incomplete, it proved difficult to craft a firm estimate, but we believe a minimum of 16,000 additional supporters are contained in that set. With caveats, we estimate a hard ceiling for ISIS supporters in the vicinity of 90,000 accounts; we could not establish a definitive upper limit. Based on anecdotal observation, we suspect the true number does not approach this level; however the metrics described in section 3.8 allow for this possibility.

All data in this paper pertains to specific ranges of time when the data was collected, from October 4 through November 27, 2014, with some seed data retrieved in September 2014. Thousands of accounts were suspended and created throughout the period of data collection. Therefore this estimate does not reflect the exact user base of ISIS at any specific moment, but rather reflects activity on a rolling basis. The user base at any given moment was likely smaller than the total estimate of 46,000.

The only way to capture a snapshot over a more condensed time frame would involve either directly accessing data, with permission, from within Twitter's own systems, or violating Twitter's terms of service regarding the speed of access of data. These options were respectively unavailable and undesirable.

When coding samples, we adopted a very conservative regimen, detailed below, which likely underestimates the amount of support for ISIS in the dataset by emphasizing overt support and excluding ambiguous classes of accounts. There are three ambiguous classes of account that should be considered when evaluating these results:

- **Covert supporters of ISIS:** Users who took medium to strong steps to conceal their support due to fear of prosecution or suspension by Twitter. Users who took only casual steps to disguise their support were generally detectable.
- **Pro-ISIS intelligence operatives:** Some users who follow accounts related to the enemies of ISIS, such as rival jihadists, would be coded as non-supporters under the conservative criteria we employed.
- **Anti-ISIS intelligence operatives:** These are accounts created to appear as ISIS supporters in order to allow ISIS's enemies to monitor its activities, which would be coded as supporters (if done effectively).

After reviewing hundreds of accounts in the set—with a focus on those that appeared ambiguous—we believe a significant number of accounts in the Demographics Dataset fall into the first two

categories. Nevertheless, most were coded as non-supporters. We did not develop a methodology to evaluate the third category, and the number of potentially relevant users remains unknown.

Several variations of our identification methodology produced extremely similar results in the top 20,000 accounts, adding to our confidence in the integrity of the sample. The final metric was more effective at lower ranges as well.

The number of supporters has certainly changed since the data was collected; we provide data on more recent changes in section 2.5.4.

1.2 The Demographics Dataset

After determining the estimated total number of ISIS-supporting accounts, we sought to describe a representative sample as completely as possible.

Because the quantity of data analyzed was too large to allow for an individual review of every single account, we had to sort ISIS supporters from non-supporters among accounts for which we had robust data.

Non-supporter accounts in the data collected included enemies of ISIS, non-ISIS jihadis, people tracking the organization's activities (such as journalists and researchers), and accounts for online services used by ISIS supporters, such as @YouTube or @Twitter.

Using metrics described in section 3.6, we sorted the 49,379 accounts for which we collected full data according to the probability that an account belonged to an overt ISIS supporter. We evaluated the performance of the metrics by coding samples from the dataset using a conservative methodology for identifying visible supporters. We also weeded 5,841 accounts according to the following criteria:

- Suspended mid-collection, 90 accounts (see section 2.5.1)
- Accounts with more than 50,000 followers (see section 2.2)

- Accounts primarily controlled by bots or apps (see sections 1.13 and 3.3)

This first round of eliminations left us with 43,538 accounts. We identified a set of 20,000 accounts, of which we estimate more than 93 percent are ISIS supporters, with a margin of error of about +/- 2.54 percent. This set of 20,000 accounts, the "Demographics Dataset," was used to produce the descriptive analysis in this section, except where explicitly noted.

1.3 Data Snapshot

Best estimate of total number of overt ISIS supporter accounts on Twitter: **46,000**

Maximum estimate of ISIS supporter accounts on Twitter: **90,000**

Number of accounts analyzed for demographics information: **20,000**

Estimated percentage of overt ISIS supporters in demographics dataset: **93.2 percent (+/- 2.54 percent)**

Period over which data was collected: **October 4 through November 27, 2014, with some seed data collected in late September 2014**

Top Locations of Accounts: **“Islamic State,” Syria, Iraq, Saudi Arabia**

Most common year accounts were created: **2014**

Most common month accounts were created: **September 2014**

Number of accounts detected using bots and deceptive spam tactics: **6,216 using bot or spam technology for some tweets; 3,301 accounts were excluded from the Demographics Dataset for primarily sending bot or spam content**

Average number of tweets per day per user: **7.3 over lifetime of account, 15.5 over last 200 tweets by user**

Average number of tweets per user (Over lifetime of the Account): **2,219**

Average number of followers: **1,004**

Smartphone usage: **69 percent Android, 30 percent iPhone, 1 percent Blackberry**

1.4 Location

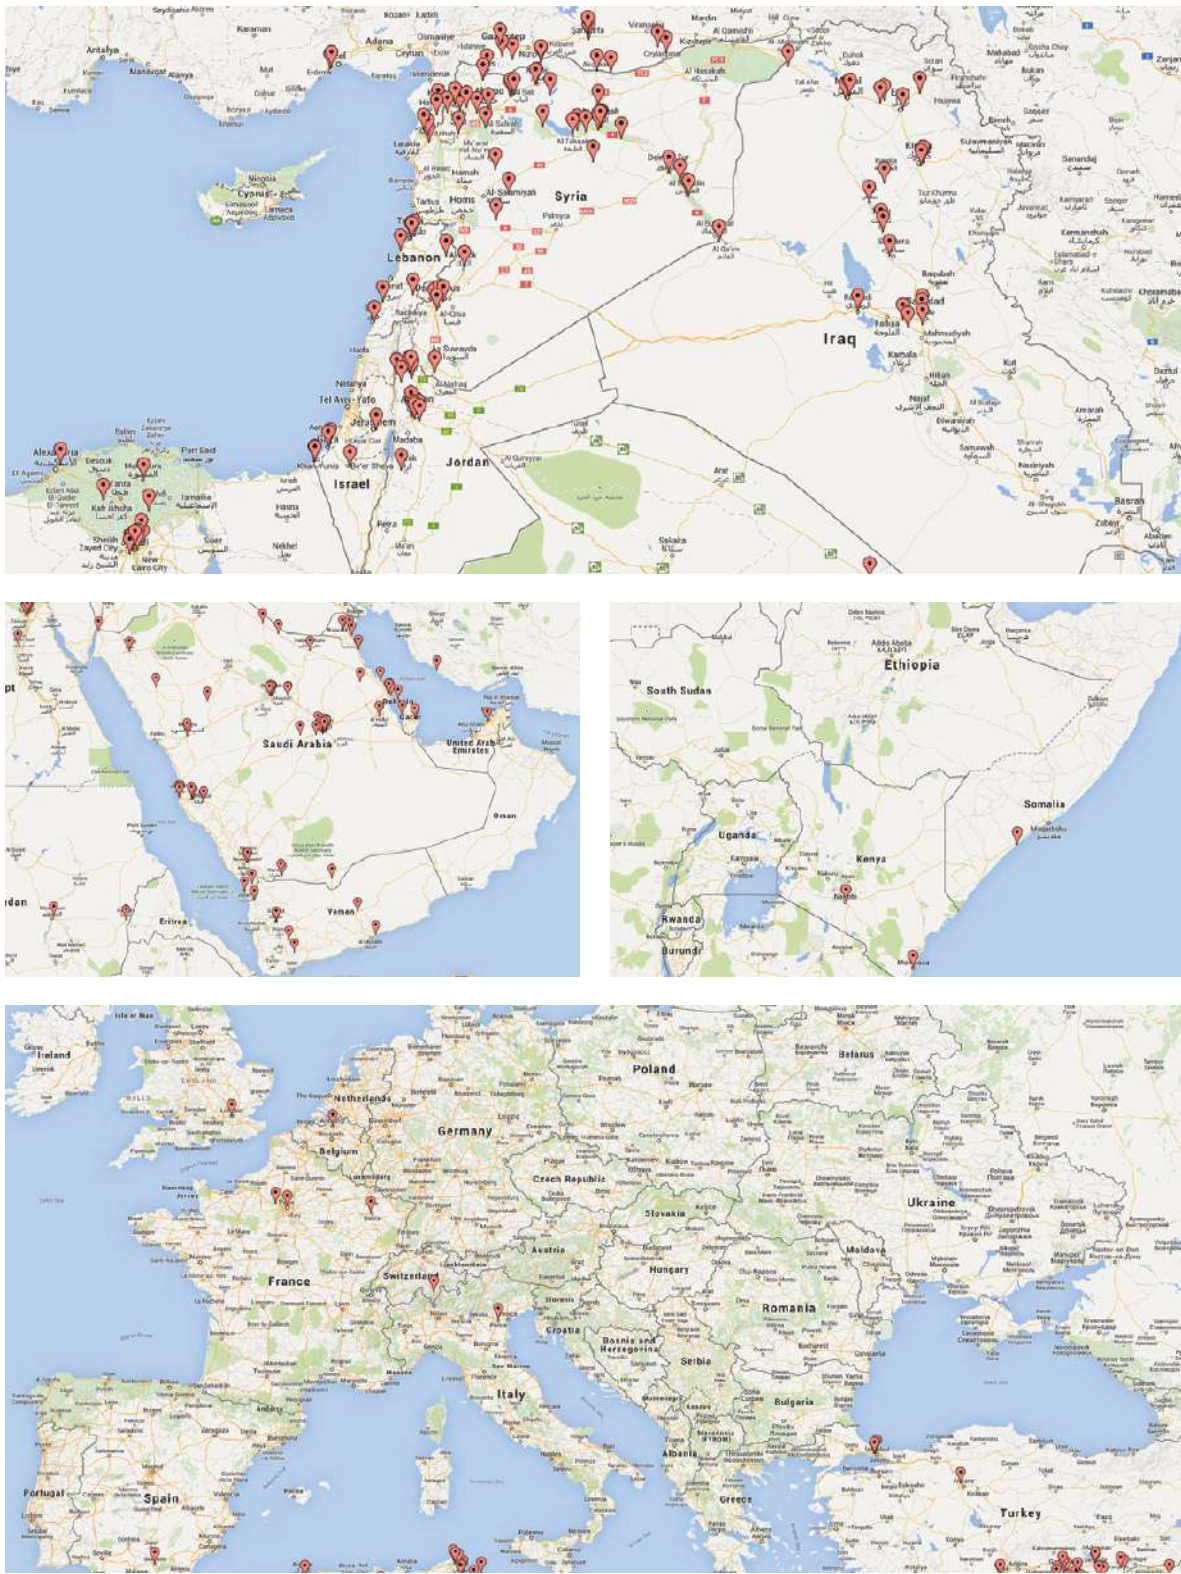


FIGURE 1: Likely ISIS supporters who sent at least one tweet out of their last 200 with location metadata enabled in the fall of 2014. Additionally, two users were located in Brazil, and one each in Indonesia and Australia.

One of the most immediately interesting and important questions about ISIS's online supporter base is where users are located.

Using open source means, the only totally reliable method of geo-locating users is to obtain coordinates provided when a user has enabled the location feature on his or her smartphone (usually resolving to a GPS signal or a cell phone tower). We analyzed only the most recent coordinate provided by each user.

Unsurprisingly, very few users in the dataset opted to enable coordinates; the number who did was surprisingly high given the operational security implications. Confirmed ISIS supporters used location services less frequently than non ISIS-supporting, typical users, but not dramatically less.

Out of the 20,000 users in the Demographics Dataset, 292 had enabled location on at least one tweet out of their last 200, or 1.5 percent. By way of comparison, a 2013 study of Twitter activity found that on a typical day in 2012, 2.02 percent of all users had location data enabled, a figure that has steadily increased over time.¹ For the entire Census Dataset of 43,538 accounts, the figure was 3 percent.

The largest cluster of location-enabled accounts (28 percent) was found in Iraq and Syria, mostly in areas either controlled or contested by ISIS. More than twice as many users reported coordinates in Syria than Iraq. The next most common location was Saudi Arabia, with 27 percent. After Syria, Iraq, and Saudi Arabia, no single country represented more than 6 percent of the total.

None of the location-enabled users were based in the United States; Western countries showed only single-digit totals (e.g., three accounts in France; two in Brazil; one in the United Kingdom; one in Australia; one in Belgium).

In January 2015, a user presenting himself as an ISIS supporter was observed using a detectable app to create tweets associated with falsified GPS coordinates for Mosul. While we cannot rule out the possibility of such tactics within the accounts examined in this paper, we did not observe evidence of their widespread use.

ISIS documents circulated on paper in Iraq, and later posted to social media, indicated that the organization's leadership was deeply concerned about the use of smartphone GPS location by its supporters and members on the ground.

In mid-December 2014, ISIS ordered members to disable GPS on their mobile devices within one month, warning that violators would have their phones confiscated and destroyed. In early January, we collected a fresh sample of tweets from a similar group of users which suggested this guidance had thus far been widely disregarded.

Because so few users (statistically speaking) had enabled location, we also analyzed location using a number of different data points provided by accounts. These locations reflect information the users chose to share, and may be deliberately misleading (in some cases demonstrably so). There were other complications as well, detailed below.

1. Kalev Leetaru, Shaowen Wang, Guofeng Cao, Anand Padmanabhan, and Eric Shook, "Mapping the global Twitter heartbeat: The geography of Twitter," *First Monday*, Vol. 18, No. 5-6, May 2013, <http://firstmonday.org/article/view/4366/3654>; Jalal Mahmud, Jeffrey Nichols, and Clemens Drews, "ACM Transactions on Intelligent Systems and Technology (TIST)," Special Section on Urban Computing archive, Volume 5 Issue 3, September 2014, <http://arxiv.org/ftp/arxiv/papers/1403/1403.2345.pdf>.

1.4.1 Inferred Location

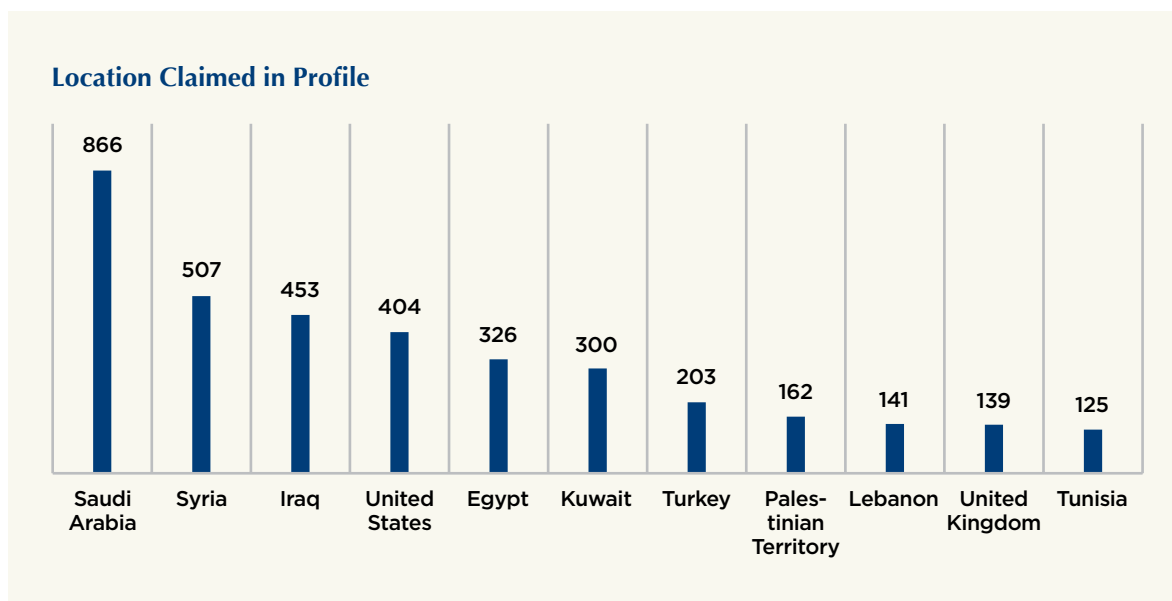


FIGURE 2: Top locations claimed by users; users listing “Islamic State” were distributed between Syria and Iraq at a ratio equivalent to the distribution seen in location-enabled tweets.

In addition to GPS coordinates, there are a number of ways to infer a user’s location based on information the user has provided. Each of these comes with a tradeoff, usually in the form of *more* data versus *higher quality* data.

Aside from location-enabled data and the content of tweets—deemed too noisy to be reliable—users can provide information about their location in the following formats:

- The “location” field on their Twitter profile. Users can enter several words in the field with no restriction on content; for example, they can offer jokes or non-relevant information.
- Time zone: Users can select whatever time zone they want. Because the selected time zone influences the time of tweets appear on a user’s timeline, there is motivation to enter it correctly, although many users do not.
- The “Bio” field in their Twitter profile, where users can write 160 characters of descriptive text on whatever subject they like, usually about themselves.

Our efforts to infer location using a combination of these factors were complicated by the erratic patterns of data entry in the fields by users, as well as the various languages in the data set. In all permutations of the analysis, locations of Iraq, Syria, and “Islamic State” were dominant.

Since locations are free-form text fields, we used a third-party algorithm to resolve entries to the country level. The list of locations that resolved to the United States was extremely noisy, including entries such as “Earth,” “everywhere,” “in the kitchen making a sandwich,” and “wherever the plane’s taking me.” However, some American cities were specified, primarily New York and Washington, D.C. Most locations could not be verified, and none of the location-enabled users were based in the United States. We are reasonably certain some ISIS supporters deceptively listed locations in the United States in order to create the appearance of a homeland threat.

Nevertheless, the location field was the only method that produced a confidence-inspiring re-

sult. Accounts that provided information in the location field resolved to 107 countries, with 960 accounts concentrated in Iraq and Syria. Users who listed “Islamic State” as their location were considered to be in either Syria or Iraq, and we assigned them to one or the other following the two-to-one distribution noted in the location-

enabled findings. Another 866 accounts claimed Saudi Arabia as a location.

With the exception of the United States, the top 20 countries correlated to regions where ISIS enjoys substantial support, such as Egypt, Tunisia, Libya, Yemen, and Gaza.²

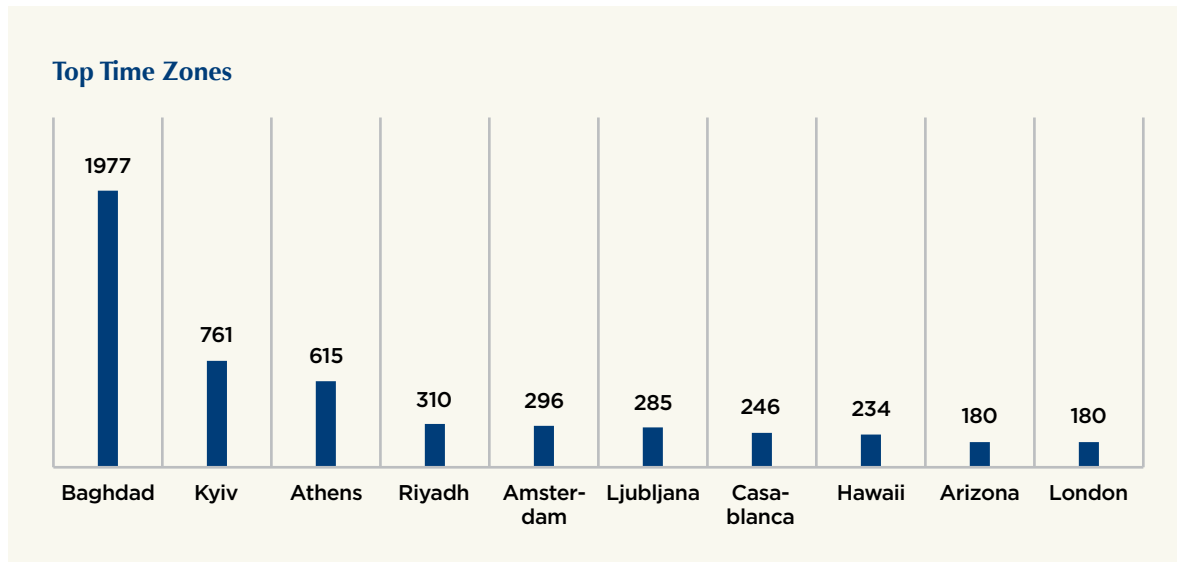


FIGURE 3: Time zones selected by users

Within the Demographics Dataset, 6,546 users opted to specify a time zone. The most frequently listed were Baghdad, 31 percent; Kyiv, 12 percent; Athens, 10 percent; and Riyadh, 5 percent. As previously noted, a notable number of users listed Arizona and Hawaii as their time zone; combined, these accounted for 6.3 percent.

Athens and Kyiv share a time zone adjacent to the Baghdad time zone and encompasses portions of Turkey and Syria, including parts of Syria where ISIS maintains a heavy presence, such as in the group’s de facto capital in Raqqa, Syria. Since Twitter does not offer a time zone selection that names a Syrian city, these entries likely reflect users in those locations.

There was a significant disparity between the stated locations and time zones of users. It is important to

note that the name of the city listed in a time zone may not reflect the actual location of the user, even if the user is in that time zone.

Location-enabled users entered both time zones and locations that did not match their actual locations. For instance, within the Census Dataset of more than 43,000, only 20 users provided all of the following: a claimed location, a time zone, and a location-enabled coordinate. All 20 of those users misrepresented their claimed location relative to the location provided by GPS.

Considerably more users designated their location as being in the Baghdad time zone than specified Baghdad as their actual time zone. It is possible that users outside that region are more willing to disclose their time zone; however, the

2. Aaron Zelin, “The Islamic State’s Model,” *Washington Post*, 28 January 2015, <http://www.washingtonpost.com/blogs/monkey-cage/wp/2015/01/28/the-islamic-states-model/>.

primary factor was likely the large number of users who listed Islamic State as their location, which may be meant as a show of support rather than a disclosure.

The vast majority of users listing Baghdad as their time zone entered a location that was metaphorical or unclear; 42 percent stated they were either in Iraq or Syria, with many listing their location simply as the Islamic State.

1.5 Languages

Twitter offers data on what language a user selected when completing his or her profile information. The choice of language is important to the process of navigating the website—for example to read menus and settings—and also dictates where such features will appear on the page.³ User language selection does not necessarily correlate to the language used in tweets. Multilingual users might choose to navigate the site in a second language, if they are adequately proficient.

Language data was available for more than 18,000 members of the Demographics Dataset, excluding some accounts which were suspended or otherwise changed status before data could be collected.

Among those users, 73 percent selected Arabic, 18 percent selected English, and 6 percent selected French, a finding that tracks to some extent with the distribution of Western foreign fighters,⁴ with an overemphasis on English that also likely reflects ISIS's target audience in the United States for inciting and harassing propaganda. No other language comprised more than 1 percent of the total.

In a 2013 study of typical Twitter users who selected a language preference, 51 percent chose

English as their preferred language, compared to 3 percent who selected Arabic. The trend of Twitter adoption by Arabic users from 2011 to 2013 suggests this number probably increased sharply in 2014.⁵

As far as content, many users tweeted in more than one language, sometimes as part of ISIS social media strategies to direct messages at external target audiences, such as when it publicizes the beheadings of Western hostages. Tweets also frequently featured a mix of languages, such as English hashtags attached to Arabic content.

3. For example, the site's design flips to the right or left depending on whether a user selects Arabic or English.

4. "Foreign Fighters in Iraq and Syria," *Radio Free Europe/Radio Liberty*, 29 January 2015, www.rferl.org/content/infographics/infographics/26584940.html.

5. Ashwin Shehagiri, "The Languages of Twitter Users," *New York Times*, 9 March 2014, http://bits.blogs.nytimes.com/2014/03/09/the-languages-of-twitter-users/?_r=0.

1.6 Display Names

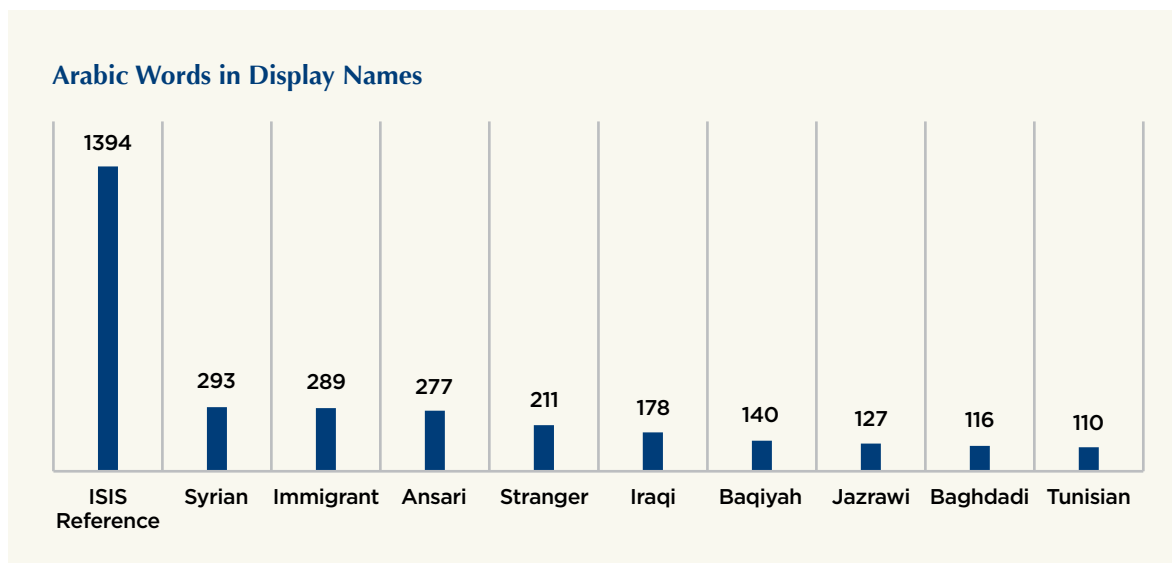


FIGURE 4: Most common Arabic words entered by users in the “display name” field

Twitter users identify themselves in two primary ways. The first is through the selection of a user handle (“@johndoe”), and the second is through the selection of a display name on the profile page, such as “John Doe,” although the display name can consist of any words or symbols up to 20 characters in length.

Despite a wide range of individual variation, some trends in the selection of display names were detected. The only powerful trend involved identification with the Islamic State in a broad sense using terms such as *Dawla* (Arabic for “state,” used as a shortened name for “Islamic State”), *baqiyah* (an ISIS slogan), *Shami* (Arabic for Syrian) and references to the caliphate. These terms were overwhelmingly more frequent than virtually any other words used in display names.

Similar to location information, user display names in the Demographics Dataset were subject to misdirection and confusion. However, some markers were commonly used to highlight foreign fighters: *muhajir*, “immigrant”, and *ghuraba’*, “strangers.” One of those two words appeared in a total of 500 user profiles. References to nationalities suggested an even larger number of foreign fighters in the set.

Within the top 20,000 users, 239 accounts were observed to use the Arabic words *umm* (mother) or *bint* (girl or daughter) in their handles or display names to indicate that they were female. In contrast, 4,536 users used the Arabic word *abu* in their handles or display names, claiming a male identity. Other gender indicators were used in names, but the complexity of navigating multiple languages and naming conventions to produce a credible result with directly comparable terms for men and women exceeded the time available for this study. However, a third-party tool relying on opaque criteria estimated approximately seven men for every one woman in the network. Additional research would provide better insight into this question.

Earlier in 2014, nearly 500 accounts identifying themselves as *umm* or *bint* were detected during an unpublished research experiment by J.M. Berger, which specifically sought to identify female supporters of ISIS. In this study, male and female social networks were observed to be segregated to some extent, often at the explicit urging of both male and female ISIS supporters, which may have also influenced the results.⁶

6. Ellie Hall, “Inside The Chilling Online World Of The Women Of ISIS,” *BuzzFeed*, 11 September 2014, <http://www.buzzfeed.com/elliethall/inside-the-online-world-of-the-women-of-isis#tj0p38MZL>.

1.7 Date of Account Creation

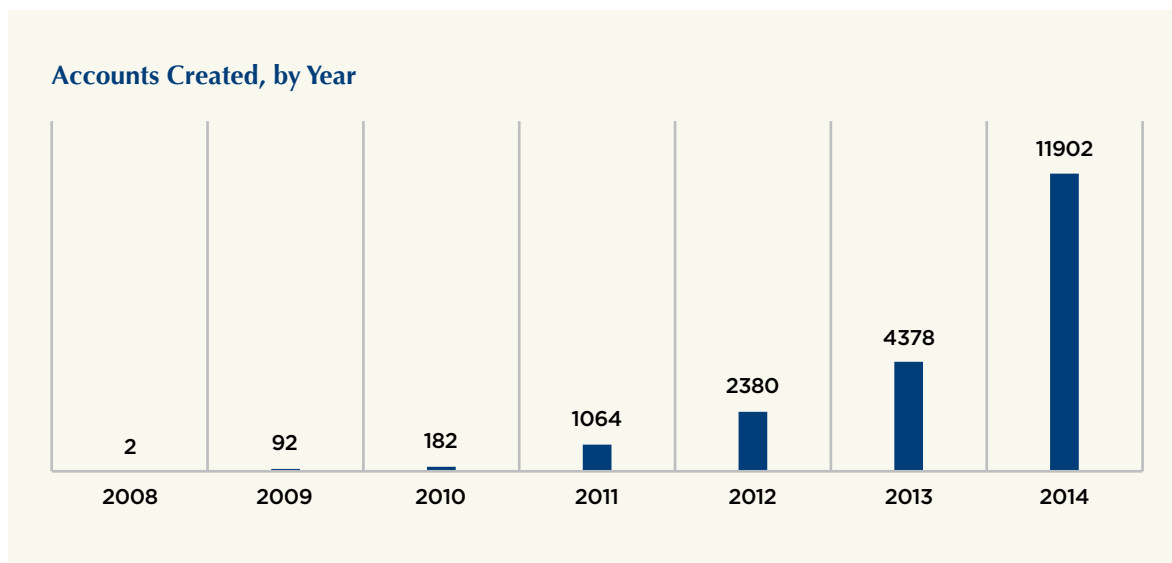


FIGURE 5: Date of account creation, by year

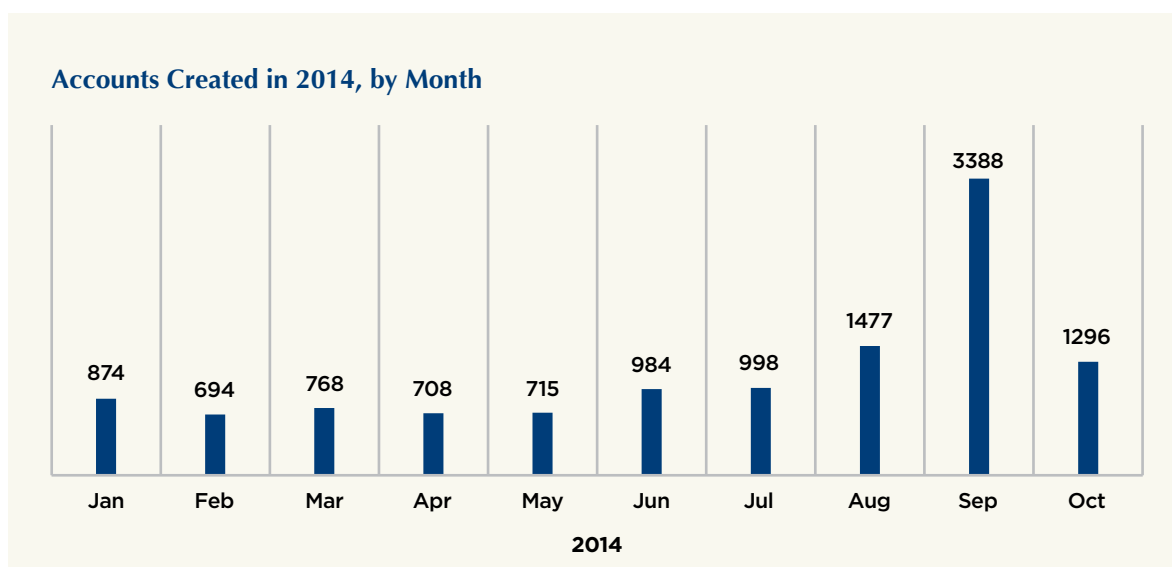


FIGURE 6: Accounts created in 2014, by month

As seen in figures 5 and 6, an extremely large number (23 percent) of ISIS-supporting accounts were created in 2014; accounts showed the most activity in September of that year. Collection of user information ended in mid-October (other

data took longer to collect); the actual total for that month would presumably have been higher had collection continued. (A differently derived comparison figure for October can be seen in section 2.5.4.)

Not coincidentally, September 2014 corresponds with the time frame during which Twitter began to aggressively suspend ISIS supporters for tweeting graphic images and videos of the beheadings of Western hostages.

We believe many—perhaps most—of these accounts were created in response to the suspensions, either to replace accounts that had been taken down, or as backup accounts to hedge against future suspensions and other steps to offset ISIS’s Twitter influence. A potentially substantial number of users may have created and recreated multiple accounts during the September 2014 time frame, which were subsequently suspended.

Only 1.3 percent of all accounts were created prior to December 31, 2010. Almost 60 percent were created in 2014. This suggests that most ISIS supporters are relatively new to Twitter, or that they created new accounts to reflect a change of interest.

It is important to account for Twitter’s overall growth in the past several years; however, the growth in ISIS supporting accounts outstripped that of the overall Twitter user population. Twitter’s user base grew by approximately 30 percent in 2013, while ISIS’s user base nearly doubled. The Twitter user base grew 20 percent in 2014;⁷ the number of ISIS supporters on Twitter nearly tripled during the same period (within the limits of the sample). This is reflective of strong growth, but also reflects anecdotal observations of increased adoption of social media by jihadist extremists starting in 2013.

The growth in ISIS’s online support base also broadly correlates with its growth on the ground in Iraq and Syria, and the course of its rift with and ultimate separation from al-Qa’ida. Some al-Qa’ida-supporting Twitter users may have created new accounts after this split to demonstrate their allegiance to ISIS.

At least some of the growth in the number of ISIS supporters is organic, although the spikes in the data clearly reflect considerations related to the suspension of accounts in September and October 2014. Arguably, the spike could have exaggerated the total estimate of supporters by as much as 20 percent. We conclude that at least some of the churn effect of accounts being created and destroyed canceled itself out.

The question of users who simultaneously maintain multiple accounts is also relevant, since some users create duplicate and backup accounts in response to being suspended. Such accounts were observed anecdotally, and most backup accounts appeared to remain relatively inactive until their predecessor accounts were suspended. Since suspensions are almost always based on the content of tweets, rather than other network characteristics, there is a strong practical argument against duplicating content in a backup account created for the express purpose of evading suspension.

Nevertheless, some users were observed using Twitter client apps that include the function (such as Hootsuite and Tweetdeck). Clear examples of this behavior were relatively rare (considerably fewer than 300 accounts in the initial, unsorted Census Dataset of more than 49,000 accounts), and users consistently employing this technique were almost entirely excluded from the Demographics Dataset consistent with the rules for bots and apps (see section 3.3).

Year Created	Number of Accounts	Percentage of total
2008	2	0.01%
2009	92	0.46%
2010	182	0.91%
2011	1064	5.32%
2012	2380	11.90%
2013	4378	21.89%
2014	11902	59.51%

7. Ben Popper, “Twitter’s user growth has stalled, but its business keeps improving,” *The Verge*, 5 February 2015, <http://www.theverge.com/2015/2/5/7987501/twitter-q4-2014-earnings>.

1.8 Date of Last Recorded Tweet

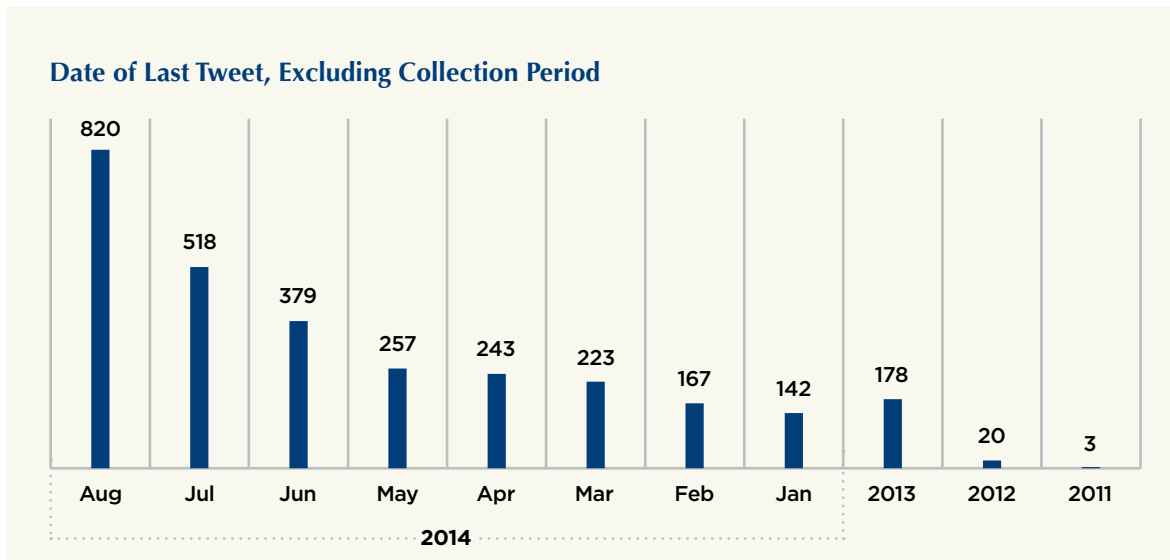


FIGURE 7: Date of last tweet, excluding the 83 percent of accounts that tweeted during the collection period

Collection of tweets took place from October through November 2014. The majority of tweets were collected in October, during which more than 83 percent of the accounts tweeted. Of the remaining accounts, 99.2 percent tweeted in 2014. That left only a fraction of 1 percent that had been inactive since 2013 or earlier, with the vast majority of those having tweeted in 2013. Given that ISIS, in its current incarnation, only came into existence in 2013, this is a logical finding.

The metrics that were used to identify ISIS supporters weight certain types of activity and may have influenced this data (section Sorting Metrics). Additionally, the criteria for seed accounts (section Starting Point) also influenced the set to favor accounts that were more active.

Other factors were weighted to avoid detecting only users with visible tweets, and we believe this data accurately points toward a high level of activity among ISIS supporters. Only 44 percent of all existing Twitter accounts display even a single visible

tweet,⁸ so by virtually any measure, ISIS-supporting Twitter users are far more active than ordinary users. A similarly derived analysis of 400 Al-Qa'ida in the Arabian Peninsula (AQAP) supporters, for instance, found that 34 percent had been silent since the start of 2014.⁹

A number of the accounts which had not tweeted since 2012 had clearly been active more recently, since some of them were branded with ISIS's name established after 2013. In some cases, we discovered in later investigation that such accounts were actively tweeting, suggesting that some users within this small subset deleted their tweets periodically (consistent with previous observations of this activity). Others simply did not tweet publicly; 30 percent of the 1,264 accounts for which no tweet date was collected were marked private. The remainder had no tweets available for analysis at the time of collection for whatever reason.

8. Yoree Koh, "Report: 44% of Twitter Accounts Have Never Sent a Tweet," *Wall Street Journal*, 11 April 2014, <http://blogs.wsj.com/digits/2014/04/11/new-data-quantifies-dearth-of-tweeters-on-twitter/>.

9. Analysis generated in January 2015 using same techniques as in this study.

1.9 Avatars

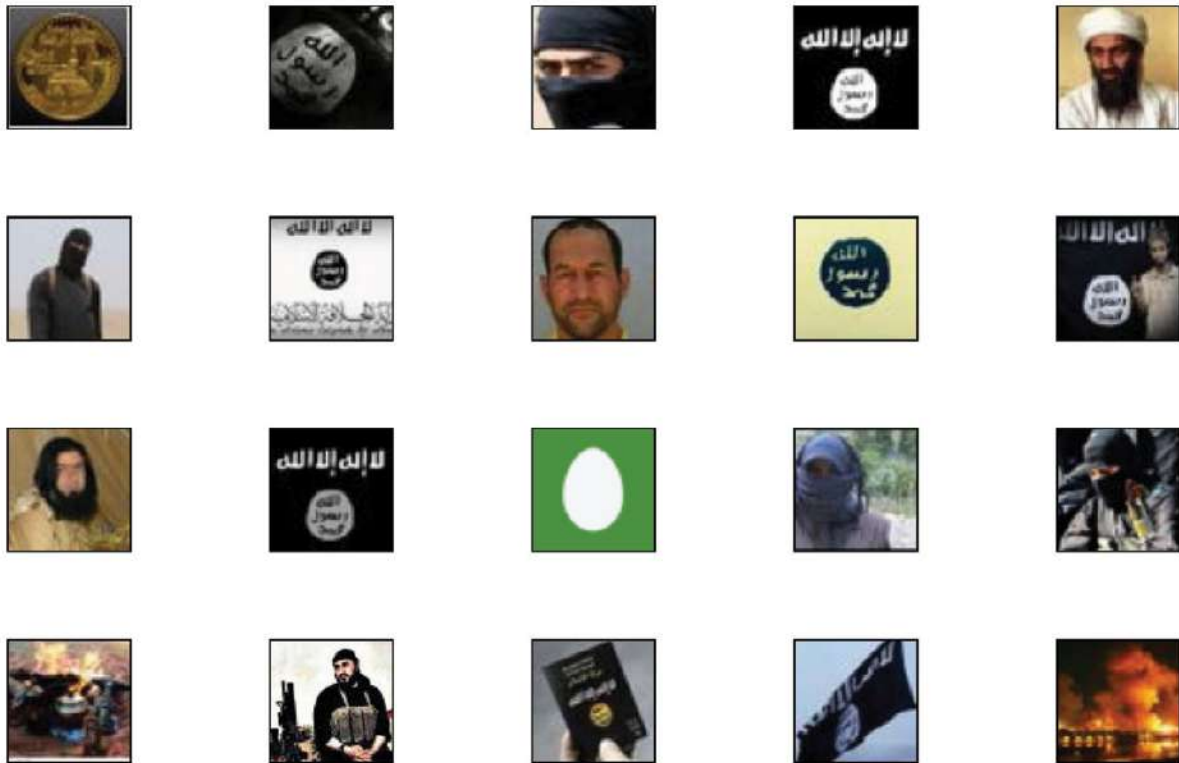


FIGURE 8: Typical Twitter profile pictures used by ISIS supporters include variations on the flag used by ISIS, images of al Qaeda founder Osama bin Laden, and prominent ISIS members and leaders including “Jihadi John” and Abu Bakr al-Baghdadi.

ISIS supporters use a wide variety of images as identity markers on Twitter. A non-exhaustive review of Twitter profile pictures within the dataset shows that these images are most prolific among users who have 150 to 1000 followers. Smaller accounts often use the default Twitter “egg,” while larger accounts sometimes employ imagery unrelated to ISIS—perhaps meant to hedge against suspension by Twitter. Within the Demographics Dataset, 6.5 percent of accounts used the default profile picture.

Images were reviewed manually rather than programmatically, although such an approach could be the basis of future research. Anecdotally, the most common profile picture was easily the iconic black and white flag used by ISIS in its official documents and propaganda, and liberally displayed in territories ISIS controls. The next most commonly-used imagery involved variations on pictures of ISIS’s leader, Abu Bakr al-Baghdadi.

Al-Qa’ida founder Osama bin Laden and al-Qa’ida in Iraq founder Abu Musab al Zarqawi were also well-represented, along with prominent ISIS members such as “Jihadi John.” Users also displayed figures less known to outsiders, such as Abu Abed Abdul Rahman al-Bilawi, who was killed while playing a key role in ISIS’s June 2014 attack on Mosul, Iraq.

Users can also select a background picture to run in a banner size across the top of their Twitter profile page. These were less consistent than profile pictures. Images often included fanciful depictions of the Islamic State, such as fighters on horseback carrying ISIS’s flag, airliners and buildings festooned with the flag, or in one case, a large sailing ship with the flag as its main sail. Other popular background images included scenes from ISIS video productions.

1.10 Top Hashtags

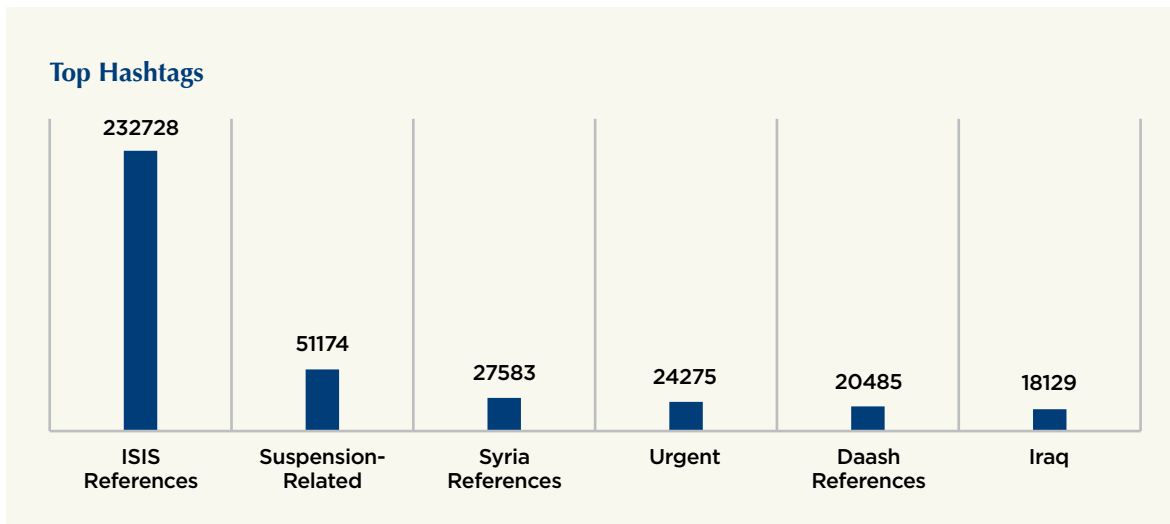


FIGURE 9: Number of tweets containing a hashtag, consolidated by theme from among the top 100 hashtags

We collected each user's 200 most recent tweets. A total of 5,384,892 tweets were analyzed, containing 100,767 unique hashtags used a total of 1,465,749 times, which is an average of once every 3.7 tweets.

At least 151,617 hashtags that included one of four most-common variations on the spelling of "Islamic State" in Arabic were detected (all of the variations could not be fully accounted for), representing 2.8 percent of all tweets.

No other hashtags even came close to that rate of usage. The next most common hashtag was the Arabic word for "urgent," which was liberally appended to news out of Syria and Iraq; this appeared in 24,275 tweets, or less than 0.5 percent of all tweets. The word "Syria" in both Arabic and English was slightly lower, with 23,769 tweets.

A breakdown of the top 100 hashtags found that 26 percent consisted of the "Islamic State" hashtag in its four most common variations. "Urgent" and Syria represented about 4 percent each.

In fourth place—representing 3 percent of the top 100 but only 1.2 percent of all hashtags—was "*Da'ish*," the

Arabic acronym for ISIS that is generally viewed as a derogatory term. The hashtag's presence reflected negative content about ISIS from non-responsive accounts in the dataset, retweets of content critical of ISIS that supporters wished to respond to, and a relatively small but notable group of users using the hashtag to send pro-ISIS messages to ISIS critics, or to reclaim the term in a positive light.

Within the top 100 hashtags, we categorized all general references to ISIS or the caliphate, all references to Syria and all references to Da'ish, as well as all hashtags pertaining to the suspension of ISIS-supporting accounts and the announcement of replacement accounts.

ISIS references represented 40 percent of the top 100 hashtags. Hashtags used in reference to Twitter suspensions were second, with 9 percent.

This figure accords with previous research by J.M. Berger based on an analysis of approximately 3,000 tweets sent from September 29 to October 1 of 2014, which found that at least 8 percent of ISIS tweets sent during that period pertained to account suspensions.¹⁰

10. J.M. Berger, "Resistible Force Meets Movable Object," *Intelwire*, 2 October 2014, <http://news.intelwire.com/2014/10/resistible-force-meets-movable-object.html>.

1.11 Top Links

Tracking the external content (such as links to propaganda) sent by ISIS Twitter accounts forms an obviously important function. However, we could not satisfactorily resolve this during the course of this study, due to a number of complications. We therefore recommend further research on the topic.

The calculation of most-linked content was complicated by a number of factors, most importantly Twitter's URL shortening practices and the removal of ISIS content by third-party Internet service providers such as YouTube.

Twitter employs URL-shortening to allow users to send links to content without exceeding the 140-character limit for a tweet. When a full-length URL (such as www.youtube.com/watch?v=h77URBgDRlc&feature=youtu.be) is entered by the user, it is shortened to take up fewer characters (such as youtu.be/h77URBgDRlc). Twitter keeps track of shortened URLs by shortening them further for internal use, rendering them all under the domain t.co (t.co/KcEfbPiWgc).

We detected 2,149,327 shortened URLs in the set of 5,384,892 tweets, or one URL for every 2.5 tweets. There were more than 1.5 million unique shortened URLs, but since multiple shortened URLs can point to the same target URL, we are certain that fewer unique web pages were linked. We were unable to resolve this question in more detail.

Expanding these URLs after the fact was problematic. Some users shortened a URL before entering it into Twitter, so expanding the t.co URL only points to another shortened URL. If Twitter deems content to be a violation of its terms of service, the t.co link will simply stop functioning.

Of the unique t.co URLs, 689 were tweeted more than 100 times, with the first-ranked URL being tweeted only 846 times. Many of the URLs were malformed, meaning the t.co URL could not be extracted accurately from the text. We attempted to expand the top 20 unique URLs to discover the content they pointed to, with the following results shown in the table below:

Shortened URL	Expanded URL
http://t.co/wdrKsuc5fd	http://retweetcom.com/ret
http://t.co/Fh1Y73OyEX	http://twitter.com/by3_S/status/519152478433345536/photo/1
Malformed URL	Expansion failed
Malformed URL	Expansion failed
http://t.co/H6uUirXBD9	http://retweetcom.com/ret
Malformed URL	Expansion failed
http://t.co/kPcWorCHG8	http://twitter.com/ilyass_4/status/519149486795272192/photo/1
http://t.co/m2azV15Yyu	http://twitter.com/Fighter_Otaibi/status/521357691491336192/photo/1
http://t.co/	Content removed due to personal identification
http://t.co/mFywYHZp	Expansion failed
http://t.co/PygkTd7N2n	http://rt10.a77mad.com
Malformed URL	Expansion failed
http://t.co/QsnbwSILwE	http://topretweet.com/
http://t.co/pbKMmJEnTH	http://twitter.com/ilyass_4/status/519149486795272192/photo/1
Malformed URL	Expansion failed
Malformed URL	Expansion failed
http://t.co/Pr6plIGs1Z	http://cutt.us/fnzQF
Malformed URL	Expansion failed
Malformed URL	Expansion failed
Malformed URL	Expansion failed

Significantly, the first- and fifth-ranked URLs for which we could resolve a target page pointed to a single page that could not be loaded at the time of writing, on a domain that could also not be loaded. The domain's name, Retweetcom.com, obviously suggests it is a Twitter app for manufacturing artificially inflated retweets, consistent with manipulative tactics observed to be used by ISIS supporters on social media, and nearly 3,000 tweets in the dataset were attributed to the site's Twitter app. The expanded URLs ranked 11th and 13th pointed to similar spam-like services.

One t.co URL pointed to another URL shortening services (cutt.us), which led to yet another URL shortening services that required a login.

Of the remaining URLs, four pointed to tweets containing photos. Three of the accounts linked had been suspended; the fourth pertained to an ill child in Syria, and not ISIS. The picture was no longer available, although the text of the tweet remained; it had been retweeted 46,000 times. Because the individual was not linked to ISIS and provided personally identifying information on his account, we have removed the URL of the linked tweet.

1.12 “Official” Accounts

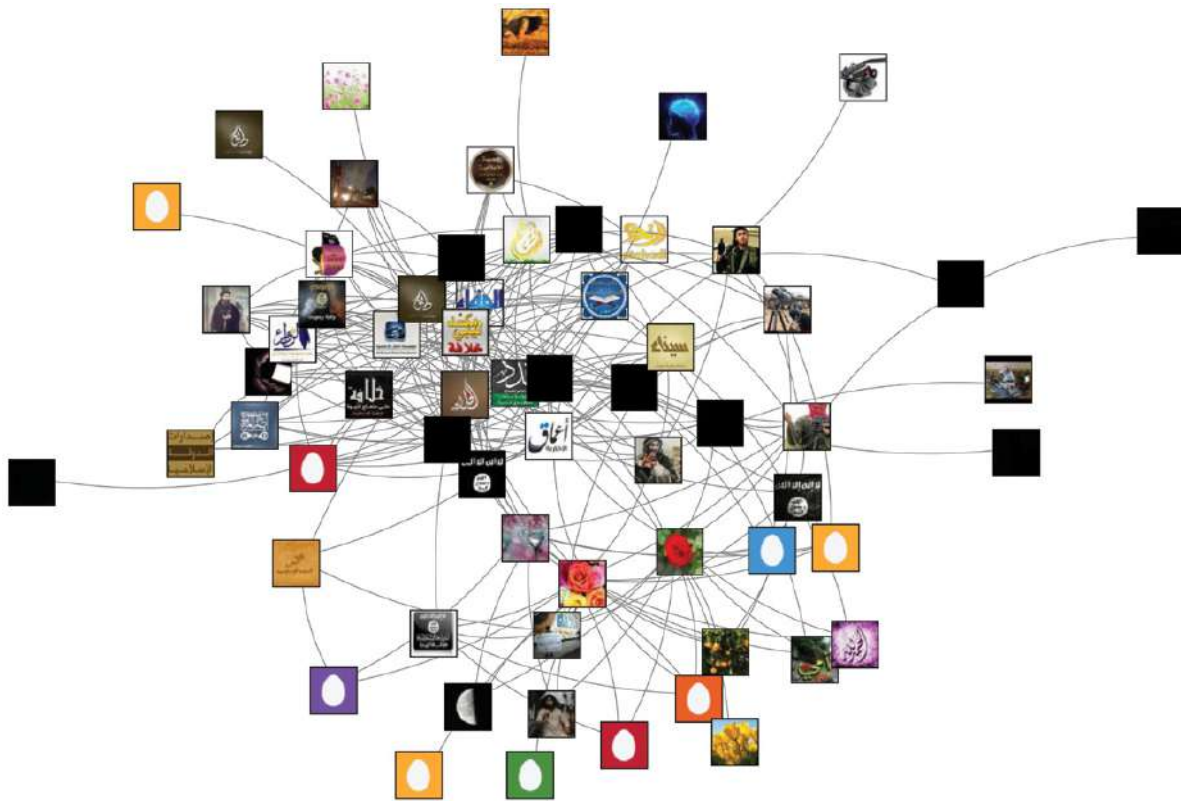


FIGURE 10: Network relationships among “official” ISIS accounts as of January 3, 2014

Twitter sporadically suspended ISIS’s primary official account throughout 2014, before taking a more aggressive stand starting in summer 2014, when it increasingly suspended most official accounts including media outlets, regional hubs, and well-known members. ISIS briefly experimented with transferring its official accounts to other social media services, where it was also met with repeated suspensions.

Although ISIS’s public, official accounts have more or less been eliminated, it has adopted coping mechanisms to maintain control over information flow on Twitter.

Specifically, its official social media operatives have established small accounts, some of which fly under the radar, while others are periodically suspended and regenerated. These users are responsible for uploading ISIS content to file-sharing and video

web sites, and then publishing links to the content. Other users (known as the *mujtabidun*, and discussed further in section 2.1) then disseminate the links more widely.

As of January 2015, ISIS had reconstituted its regional accounts with strong privacy settings, allowing only a small group of known ISIS supporters to follow the accounts and read their tweets. The content of the tweets—primarily news releases, videos and photos from ISIS’s various provinces—are then disseminated by a number of other smaller accounts using hashtags. After the initial dissemination, the content is more widely distributed, but at significantly reduced levels from early 2014.

As of December 28, 2014, we had detected 79 such “official” accounts, mostly through specific investigation and manual search, rather than via the Cen-

sus data collection process; certain high-scoring accounts from the Census Dataset were added to the list after examination. Of the 79 accounts, 24 had designated their tweets and follower lists as private. We were nevertheless able to gather some information through various analytical approaches.

The public accounts had an average of about 6,437 followers each, although fresh suspensions reduced that number within just a few days. These accounts tweeted an average of 11 times per day over their

lifetime, and 11.8 times per day within each account's last 200 tweets.

The private accounts averaged 150 followers each, and many used the default Twitter “egg” profile picture in order to maintain a lower profile. The regional accounts were mostly marked private, although some were seen to switch to public tweets at times, and they were connected to other private accounts that appear to be an essential part of the functioning of the semi-official account network.

1.13 Bots and Apps

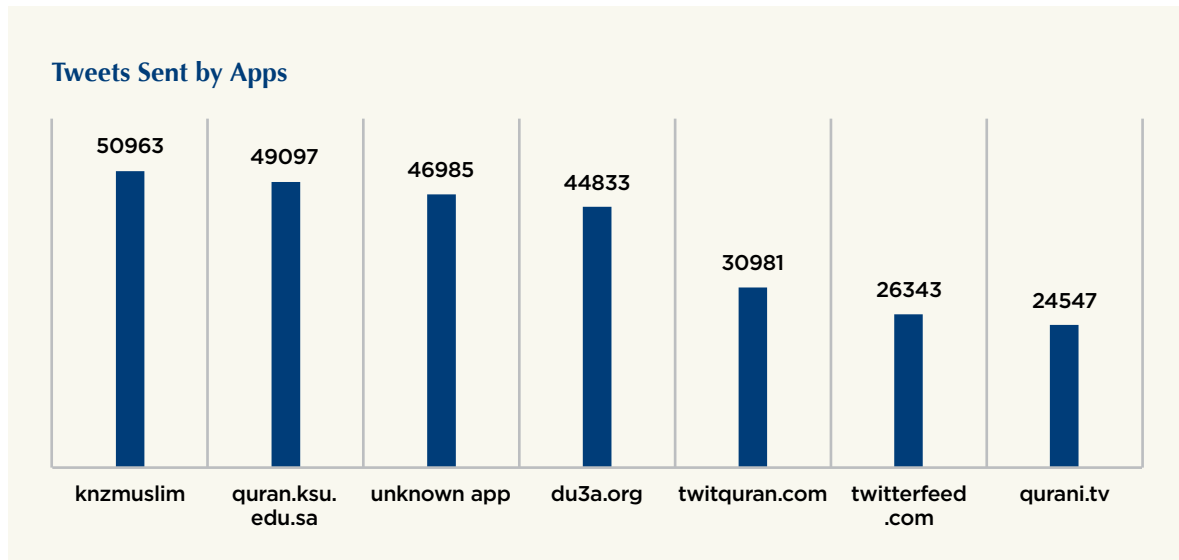


FIGURE 11: Top non-client apps used to send tweets by members of the Demographics Dataset, by number of tweets sent

ISIS supporters use a wide variety of bots and apps for many different purposes.

Bots and apps are small pieces of computer software or third-party services designed to promote content from a Twitter account automatically, without a human being manually sending tweets. A wide variety of bot and app techniques were observed in the collected data. Spam services sell tweets and retweets of selected content. A user purchases tweets from the seller, then the seller

sends out the content. These services also typically rely on bots and apps to do their work.

Some apps, such as Hootsuite and Tweetbot, are simply Twitter clients that allow users to tweet and follow others, and may not include manipulative features. We eliminated the most popular clients from the list and focused on non-client apps.

Some apps are apparently devotional in nature, tweeting prayers, religious aphorisms, and content

from the Quran, although they may also serve as identity markers or fulfill some kind of signaling function. These apps, such as knzmuslim and du3a, can produce staggering numbers of tweets per day. Knzmuslim was clocked at more than one million tweets per day, and around 1,000 tweets per minute, at one point in early January. The content they post does not overtly pertain to ISIS. In addition to their wide popularity both within and outside of ISIS circles, these apps introduce noise into social networks and their use may be intended to impede analysis.

Some apps, such as BufferApp, are used to schedule tweets to be sent at a particular time, and do not necessarily denote manipulative behavior. A number of similar commercial social media marketing tools, including some used by ordinary businesses and brands, were detected.

Other apps are intended to disseminate ISIS propaganda at a pace and volume that enables their wider distribution. The most successful of these was known as the “Dawn of Good Tidings.” In mid-2014, thousands of accounts signed up for the app, which was endorsed by top ISIS online personalities. At its peak, it sent tens of thousands of tweets per day. The app was terminated by Twitter in June 2014, silencing thousands of ISIS-supporting accounts overnight.¹¹

In the wake of that setback, ISIS supporters have responded by creating a large number of bots in small clusters, with each cluster using a different service to post tweets of the propaganda and hashtags it wishes to promote. If one “family” of bots is suspended, there are still many others that will continue to tweet. Thousands of such accounts were detected in the course of this analysis. Many from this new generation of bots were constructed using popular third party automation services such as IFTT (If This, Then That), which Twitter is unlikely to shut down since it is much more commonly used for innocuous purposes by ordinary users.

Because a deceptive app can function alongside a human operator of a Twitter account (sending tweets automatically while still allowing the user to tweet normally), we did not filter out all of the apps and bots we detected. In some cases, the app did not represent the majority of the content tweeted by the account. In other cases, apps were determined to be legitimate Twitter clients, rather than manipulation devices. We eliminated some known bot and spam providers, as well as accounts whose tweets were identical to tweets posted by other users in large volumes. In some cases, we evaluated the percentage of a user’s tweets that were sent by the app as opposed to other clients. We also reviewed technical signatures related to bot activity.

In the overall Census Dataset, around 400 non-client apps were detected to be in use among more than 6,000 accounts. Within the 5.4 million Demographics Dataset tweets analyzed, hundreds of additional bots and apps were also detected operating at lower volumes, enough to suggest that perhaps 20 percent or more of all tweets in the set were created using bots or apps.

Any given Twitter account can include tweets from both an app and a human user. We eliminated just over 3,000 accounts from the Census Dataset for very high levels of bot and spam activity, prior to creating the Demographics Dataset (section Bot and Spam Detection). In early 2015, we checked activity within the Demographics Dataset and discovered about 400 accounts that met the criteria used to exclude bots in the first collection. These accounts had changed their patterns of activity since the original collection.

11. J.M. Berger, “How ISIS Games Twitter,” *The Atlantic*, 16 June 2014, <http://www.theatlantic.com/international/archive/2014/06/isis-iraq-twitter-social-media-strategy/372856/>.

1.14 Smartphones

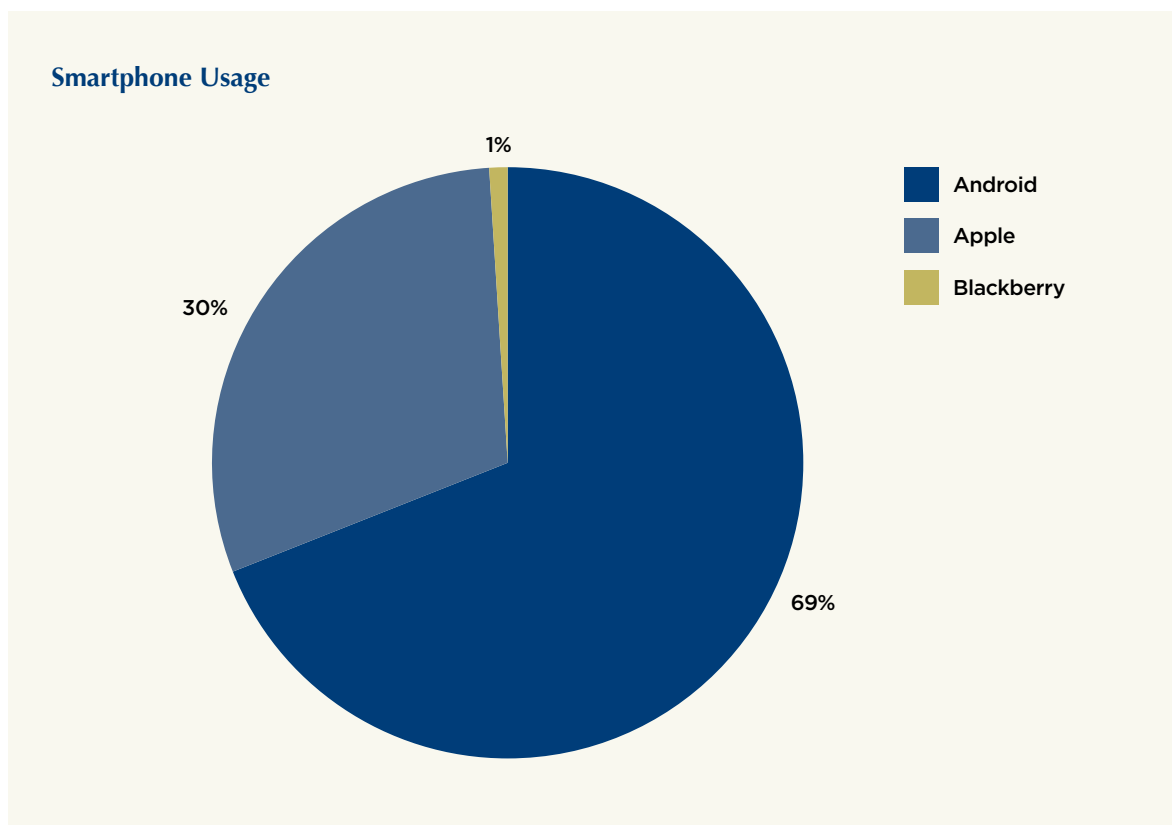


FIGURE 12: Smartphone usage among ISIS supporters, according to primary app used for tweeting

Each tweet collected specified what type of Twitter client was used to send the tweet, meaning whether it was sent from the Twitter web site or a smartphone client. For each user, we collected the name and download link for the client they employed most frequently. We broke these down according to whether the download link pointed to the Google Play store, the Apple store, or Blackberry.com. A number of other mobile clients were used, far smaller numbers, such as an app exclusive to Lenovo phones.

Among users of the three most popular phone types, 69 percent had downloaded a Twitter client from the Google Play store or Google.com. Another 30 percent used a client downloaded from the Apple iTunes store, and about 1 percent had downloaded a client from Blackberry.com.

In mid-December, ISIS announced it would ban iPhone products within its territory due to se-

curity concerns. In early February, we collected data on a set of 10,000 likely ISIS-supporting accounts using a similar methodology to the overall study, and found only a 1 percent drop in the use of iPhones.

2.

Social Media Metrics

In addition to information that reflects user preferences and location, the most basic Twitter metrics—the number of followers and following, and patterns of tweeting—was analyzed to determine the profile of a typical ISIS supporter account, and understand some characteristics of the overall network and its potential reach.

2.1 Tweeting patterns

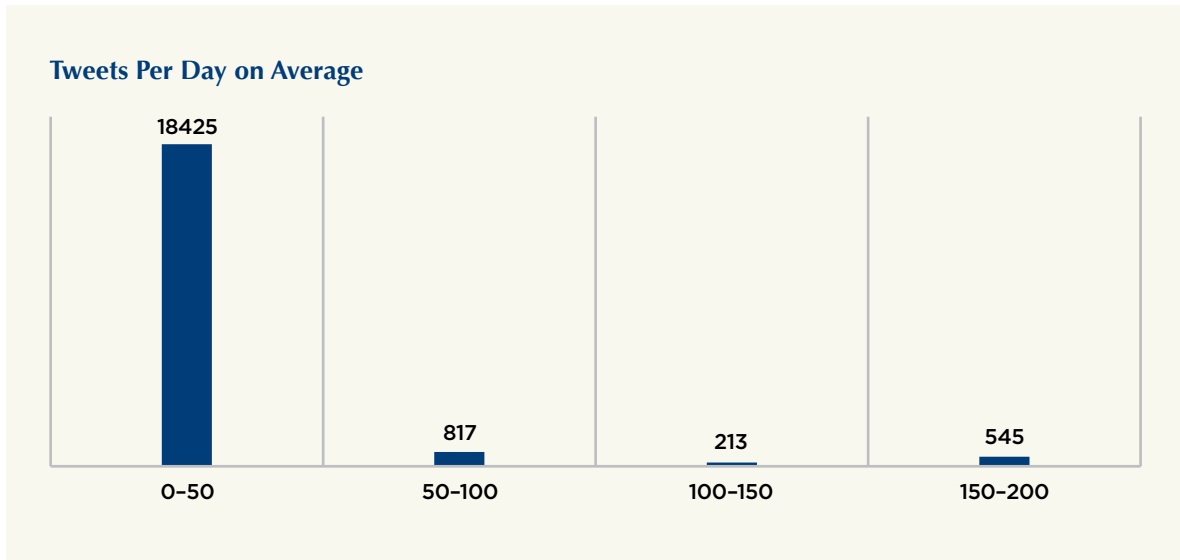


FIGURE 13: Tweets per day, calculated from each user's 200 most recent tweets

The nature of the process we used to evaluate tweets, while important to developing the metrics used to evaluate the accounts, limited evaluation of some aspects of a user's activity (especially for prolific users who tweeted at a high pace).

Therefore we evaluated tweeting patterns using two sets of data:

- Data collected for the study that analyzed activity based on the user's most recent 200 tweets. This helped us identify users who might not tweet every day or week, but who are very active when they are online.
- Data based on more than 18,000 accounts active in late December under the same usernames were identified in the initial collection of data. These figures reflected activity over the lifetime of the account.

Over the lifetime of their accounts, about 69 percent of ISIS supporters sent fewer than five tweets per day on average, with 40 percent sending less

than one tweet per day on average. Only 2.4 percent of accounts tweeted more than 50 times per day on average. A typical account tweeted 7.3 times per day, for an average daily output of 133,422 tweets per day from all members. This total output figure comes with several caveats.

These tweet-per-day averages allow for long periods during which users do not tweet. For example, a user could open an account and tweet 70 times on Monday, then stop tweeting for the rest of the week. His average would be 10 tweets per day for the week.

Overall, ISIS supporters were much more active than the average Twitter user. In the December sample, 35 percent of users had tweeted within the 24 hours preceding the collection of data, and another 16 percent had tweeted within the preceding seven days. Within the year preceding collection, 98 percent of users had tweeted at least once, considerably higher than the 67 percent of all active Twitter users who had tweeted

within the previous year. 62 percent of ISIS supporters had tweeted within 30 days of collection, compared to just 13 percent of all Twitter users.¹²

The lifetime averages, while useful, do not adequately portray user activity patterns. The most active accounts are also the most likely to be suspended (as discussed in section 2.5.1). Because of this, the December sample omitted some of the most active accounts from analysis. Furthermore, some users whose past activity was measured at a high rate had subsequently deleted all their tweets. Both of these factors shaped the late December dataset by removing data on the most prolific users.

As discussed elsewhere in this report, ISIS supporters have been observed to tweet repeatedly in short bursts in order to widely disseminate important content. Insight into this activity can be found in the original collection, which analyzed the most recent 200 tweets from each account and calculated tweets per day based on the time of the earliest collected tweet and the latest collected tweet.

This approach allowed us to detect users who tweeted in prolonged bursts, even if they later went silent for a period of time. Since only 200 tweets were collected for each user, the maximum value recorded in the dataset was 200 tweets per day, even though some users were observed to tweet more on specific days.

An overwhelming majority of users (92 percent) tweeted less than 50 times per day, including 500 accounts whose tweets were marked private and for whom no tweets were collected. However, 1,575 users tweeted more than 50 times per day on average, with 545 tweeting more than 150 times per day.

These prolific users—referred to in ISIS social media strategy documents as the *mujtahidun* (in-

dustrious ones)—form the highly engaged core of ISIS’s social media machine. These users may not tweet every day, but when they do, they tweet a lot of content in a very short amount of time.

This activity, more than any other, drives the success of ISIS’s efforts to promulgate its message on social media. Short, prolonged bursts of activity cause hashtags to trend, resulting in third-party aggregation and insertion of tweeted content into search results. Prior to the start of Twitter’s aggressive account suspensions, highly organized activity among the *mujtahidun*—who at one point we may have numbered as many as 3,000, including bots—allowed ISIS to dominate certain hashtags and project its material outside of its own social network to harass and intimidate outsiders, as well as to attract potential recruits (this is discussed at more length in section 4).

Within the entire Demographics Dataset, the average user tweeted 15.6 times per day when measured by their most recent 200 tweets, twice as high as the figure calculated over the lifetime of the accounts. Because of the 200-tweet cap on data, this estimate is certainly lower than the reality.

Since this figure reflects the most recent tweets of each user, it is capped at 200 tweets per day, is not delimited by date, and it cannot be used to extrapolate total daily volume; it does, however, clearly indicate that when users are online, they are extremely active.

12. David Murphy, “44 Percent of Twitter Accounts Have Never Tweeted,” *PC Magazine*, 13 April 2014, <http://www.pcmag.com/article2/0,2817,2456489,00.asp>.

2.2 Number of Followers

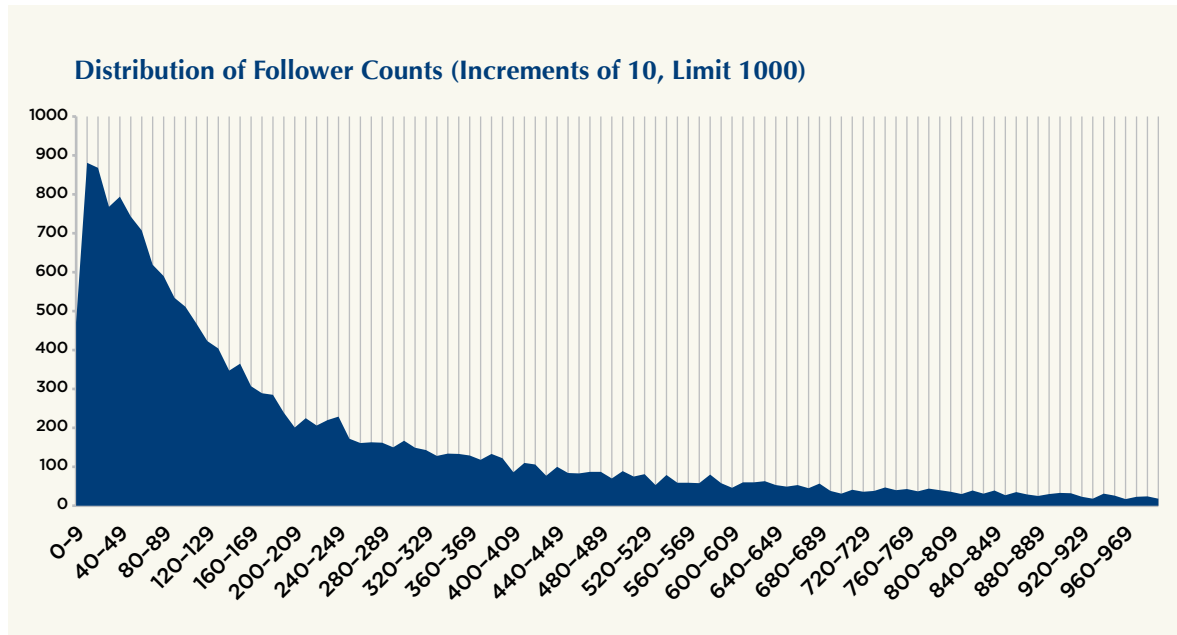


FIGURE 14: Number of accounts, ranked by the number of followers, in increments of 10, with an upper limit of 1,000

Followers counts demonstrate how limited the reach of the ISIS social network truly is, but also how it outperforms the averages.

A large majority of ISIS supporters on Twitter—73 percent—had fewer than 500 followers each. Only 4 percent had more than 5,000 followers. While these figures are very low for on-line influencers such as celebrities or mainstream politicians—who can have follower counts in the millions—they are very high relative to the average user.

The mean number of followers among ISIS supporters was 1,004, relative to 208 for the average Twitter user. The median was 177 followers, compared to 61 for the average active Twitter user, which put the typical user in the 80th percentile or higher when compared to all Twitter users.¹³ ISIS supporters outperformed typical users across the board (see figure 15), often by multiples.

The available average user figures date back to 2013 and includes a great deal of noise, thus the comparisons here are somewhat uneven. A better comparison might be to users who tweet at the same pace on a common topic, a set that excludes bots and data that was collected concurrently. Nonetheless we can confidently state that the typical ISIS supporter has more followers than the typical Twitter user. We believe future research on a number of comparable groups would be extremely illuminating.

13. Jon Bruner, "Tweets Loud and Quiet," *Radar*, 18 December 2013, <http://radar.oreilly.com/2013/12/tweets-loud-and-quiet.html>; Steven Hughes, "15 Twitter Facts and Figures for 2014 You Need to Know," *JeffBullas.com*, 16 April 2014, <http://www.jeffbullas.com/2014/04/16/15-twitter-facts-and-figures-for-2014-you-need-to-know/>.

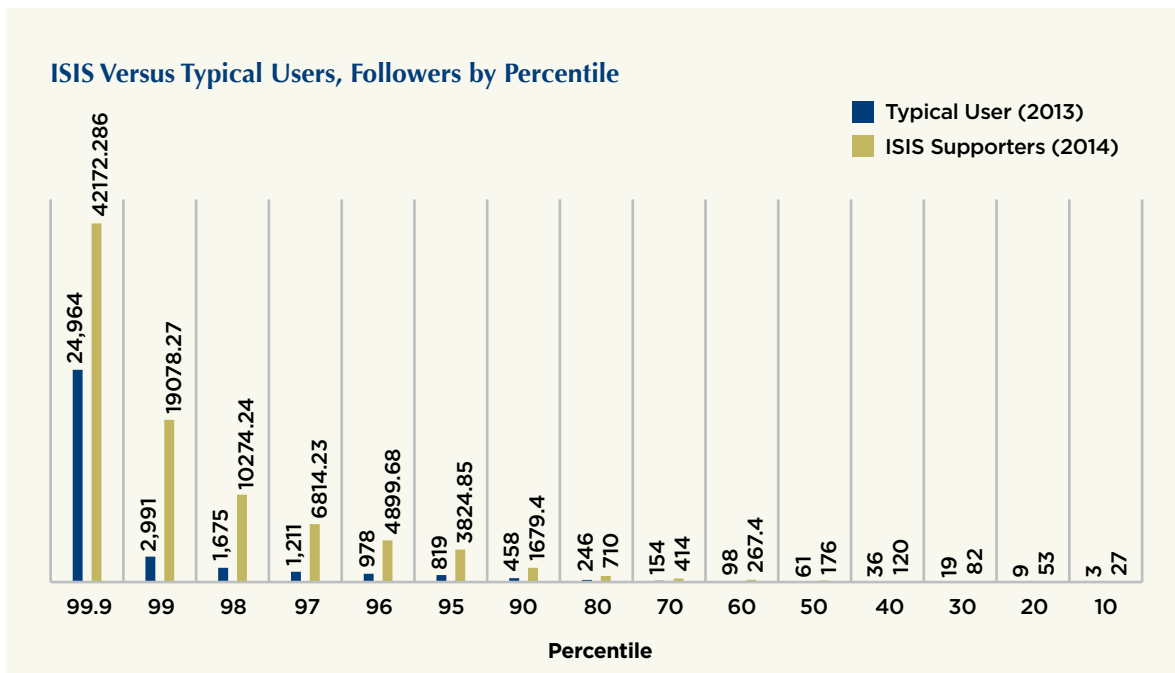


FIGURE 15: Number of followers by percentile, typical Twitter account versus ISIS supporting account

Despite its extensive size compared to average Twitter users, the ISIS follower base pragmatically represents a very small number of people, and its network is very internally focused. About 44 percent of all followers were other users within the Census Dataset, although the rapid changes in the network due to suspensions and the creation of new accounts clouded this figure. Internal focus was also one of the criteria used to identify the Demographics Set. While this raised the average in-network follower count, the criteria was chosen precisely because it correlated to ISIS supporters.

Importantly, no overt ISIS supporter had more than 50,000 followers. While some came close, few even approached that number. We could have significantly increased the accuracy of the metrics by cutting off accounts with more than 20,000 followers, but we believed this would eliminate too many responsive accounts.

Even at their most popular levels, top ISIS influencers command an audience that is fractional compared to celebrities or prominent U.S. government officials, such as President Obama (54.6 million followers), Vice President Joe Biden (735,000),

or Secretary of State John Kerry (400,000). While highly active and committed, ISIS supporters are an insignificant speck in the overall sea of Twitter's active monthly user base of 284 million.¹⁴

We do not have adequate data to fully assess the impact of previous Twitter suspensions on the user base, but it is useful to note that some ISIS supporters and official ISIS accounts racked up tens of thousands of followers at earlier points in 2013 and 2014 prior to being suspended. Suspensions tend to target larger accounts that have more followers, as well as accounts that are more active (see section 2.5 for further discussion of this issue). Even in the absence of suspension pressure, it had been rare for an ISIS supporter to register as many as 100,000 followers.

Despite the small number of followers that any given ISIS account could boast of, ISIS supporters are still highly effective at getting their message out. They employ a variety of techniques, including repeated tweets the same content by the same user within a short period of time, and tweeting the same content by many users within a short period of time, as noted in section 2.1.

14. "About Twitter Inc.," Twitter.com, 14 February 2015, <https://about.twitter.com/company>.

2.3 Number of Accounts Followed

In terms of data analysis, the number of accounts followed is somewhat less useful than the number of followers. Both figures are subject to manipulation (for instance, some ISIS accounts have been observed to have purchased followers from social marketing firms). The number of accounts that a user follows, however, is dictated by both personal preferences and specific strategies promoted by ISIS's social media team. For example, the team has issued documents recommending that users follow many accounts followed by popular Arabic preachers in order to muddy outside analysis.

A commanding majority of users—90.7 percent—followed less than 1,000 accounts each. Another 8.4 percent followed 1,000 to 2,000 accounts. Less than 1 percent followed more than 2,000 accounts.

Within the group following fewer than 1,000 users, the bias was toward smaller numbers: 81 percent followed fewer than 500 accounts, 45 percent followed fewer than 200, and 23 percent followed fewer than 100. Taken together with the follower count analysis, this again emphasizes the relatively small and cellular nature of key ISIS supporter social networks, and emphasizes the importance of its coordinated posting activity as a component of its messaging success.

Nevertheless, these figures again compare strongly to overall Twitter usage, in which many accounts sit mostly idle. The average Twitter user in 2012 reportedly followed 108 accounts, compared to a mean average of 418 among ISIS supporters (median 257).¹⁵

Two thirds of accounts followed by the average user were also within the network (meaning the Census Dataset). This is to some extent an artifact of the sorting process, which intentionally emphasized users who were more focused on the internal network.

The most followed account in the Demographics Dataset was a news source covering ISIS. While it did not overtly brand itself as an ISIS-supporting outlet, it can be considered as such based on its content. This account had only 5,462 followers when it first logged into the system, but our data showed that it added followers very rapidly (see Appendix A). At the time of this writing, the account had just surpassed 50,000 followers, making it an extreme outlier.

The second most followed account was linked to a user believed to be a Palestinian ISIS supporter with strong ties to its social media network. The account had been suspended as of the time of writing.

2.4 Follower/Following Ratio

The ratio of an account's followers to the number of accounts it is following (sometimes referred to as "friends") can act as a barometer for the influence of a Twitter account.¹⁶

We suspect the unique qualities of the ISIS dataset—including manipulative social media practices and the campaign of account suspensions—renders this measure less useful than it might otherwise be.

ISIS strategy documents urge supporters to follow large numbers of users selected semi-randomly as a hedge against account suspension; there was no indication this worked as intended, nor that the practice was widespread. In other cases, accounts using the names of certain popular ISIS supporters were observed to accrue massive numbers of followers almost instantly, thanks to a combination of bots and third-party spam services that offer to sell followers. The most egregious of these appeared briefly during the collection period and accrued almost 100,000 followers instantly. The account was suspended just as quickly, before the collection process captured a picture. Other ISIS users noted that the account was fake.

15. "An Exhaustive Study of Twitter Users Across the World," *Beevolve*, 12 October 2012, <http://www.beevolve.com/twitter-statistics/>.

16. M.G. Siegler, "Twitter's Golden Ratio (That No One Likes To Talk About)," *Tech Crunch*, 26 August 2009, <http://techcrunch.com/2009/08/26/twitters-golden-ratio-that-no-one-likes-to-talk-about/>.

All that said, the ratio of following to followers for the average ISIS supporter was 66 percent, indicating that the typical account followed more than it was followed by. Almost half of all ISIS supporter accounts (48.5 percent) followed more than twice the number of accounts that followed them. Slight-

ly more than 25 percent of all accounts were followed by more than twice as many accounts as they followed, some with very high multiples (reflecting in part, popular information sources which broadcast but do not follow). The remainder fell into a range of followers-to-friends of 0.5 to 2.0.

2.5 Suspensions

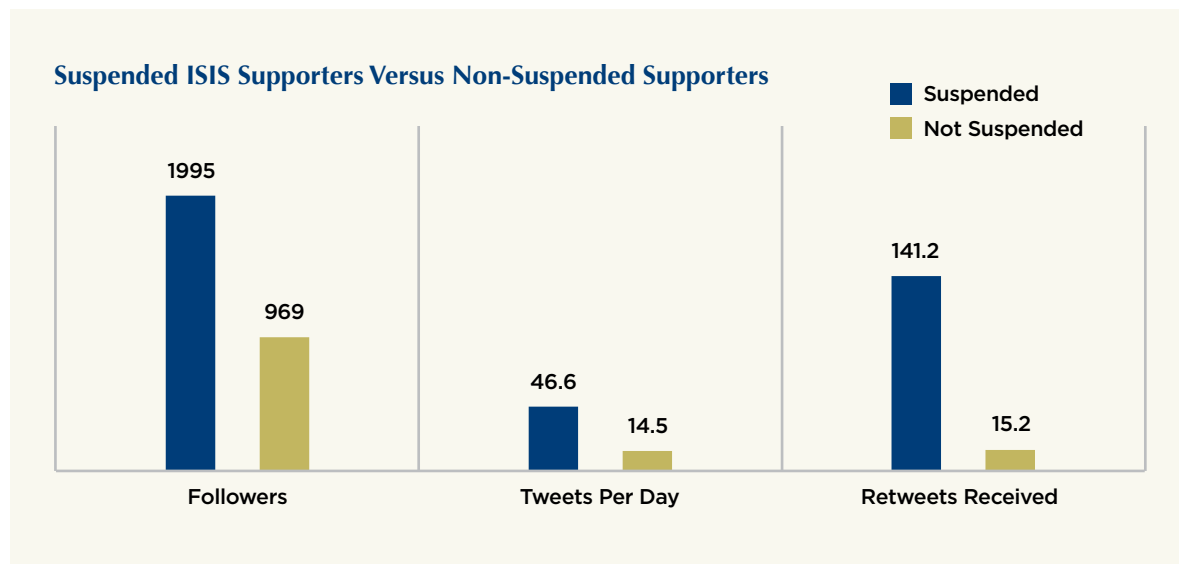


FIGURE 16: Performance of suspended ISIS supporters compared to non-suspended supporters

Analyzing the effects of Twitter suspensions on ISIS supporters is an extremely complex task for two reasons:

- No one has until now adequately publicly described the size of the community of ISIS supporters on Twitter, nor defined an adequately large sample group for study.
- Suspensions inherently change the composition of a comparison set. As a result, the manner in which comparison sets are derived is crucial to a credible assessment.

The methodology described in section 3 can provide a starting point for additional, rigorous research on this important topic by defining a method to develop clearly comparable datasets over time. As such, the data collected herein is more of a starting

point for future research, rather than a basis for a fully formed argument.

While this analysis was not designed to produce a definitive endpoint for analyzing the effects of suspensions, there are a substantial number of key preliminary data points.

2.5.1 Number of observed suspensions

790 accounts were suspended between the collection of user information in September through October and early January. Out of these, 678 were in the Demographics Dataset, with another 112 in the full collection (an expected finding given that ISIS supporters were by design more concentrated in the Demographics Dataset). Another 92 accounts were detected

as having been suspended during the collection process, and were thus omitted from the database.

The 678 suspended accounts in the Demographics Dataset had an average of 1,995 followers each, compared to an average of 969 followers for users who were not suspended. They tweeted an average of 46.6 times per day and received 141.2 retweets (in their 200 most recent tweets) compared to 14.5 tweets per day and 15.2 retweets received for users who were not suspended.

This data is unambiguous: Users who were suspended were far more active and the content they generated was far more likely to become viral than users who were not suspended. No single factor appeared to be more important than any other.

While some accounts showed little sign of activity or virality in the first collected sample, we did not monitor their activity in between the two collections, and they may have become more active prior to suspension. Had that data been available, we expect it would show that the disparity between suspended and non-suspended accounts was even higher than recorded.

In addition to the accounts for which we had collected full data, we were also able to observe a very large number of suspensions in the Level 2 set of approximately 2 million accounts, consisting of nearly 18,000 accounts. Based on the structure of the collection process, we believe a significant number of these suspensions were ISIS-related, but we could not make a definitive conclusion with the data at hand.

2.5.2 Suspensions of accounts created in September and October 2014

Significantly, 57 percent of the suspended accounts in the Demographics Dataset had been created in August, September, and October of 2014, again pointing to the possibility of repeated suspensions of single users during the period.

Our process did not fully account for suspensions of accounts created during the collection period, such as a user who was suspended, returned with a new account, and was immediately suspended again.

We speculate that this number is considerably higher than is reflected our findings and likely accounts for many of the 18,000 suspensions for which we could not obtain complete data. While our tracking of the 20,000 accounts in the Demographics Dataset only found about 678 suspensions as of January 2015, we created Twitter lists to monitor the content of the dataset during the same period, and saw many complaints about suspended accounts as well as announcements of returned accounts.

In short, the original network suffered only a 3.4 percent loss in membership between the first collection and January 2015, and the subsequent suspensions appear to have targeted new and returning accounts. The Level 2 suspension data could also point to such activity. The pace of suspensions appears to have maintained the network at a reduced level, rather than shrinking it significantly.

2.5.3 Performance of accounts that were not suspended

All of the figures above pertain to the data collected in September through December 2014. We also compared certain performance elements of non-suspended accounts collected in the fall of 2014 to their performance in January 2015.

We could not complete a full iteration generating a new and more comparable set in the time available. There were a number of problems with using only the non-suspended accounts as a comparison set:

- The original set is not properly considered a pre-suspension set, since a very aggressive regimen of suspensions started shortly before our data collection began and continued throughout this process.

- New users who should properly be accounted for in evaluating the network's performance were omitted from the second collection.
- Terrorists attacked the office of the Charlie Hebdo magazine in Paris in the midst of the January collection, resulting in a spike in coordinated tweeting activity from ISIS supporters being driven by influential users, as well as a wave of additional suspensions and name changes. We were unable to quantify this activity or to estimate its impact beyond simply noting it was observed.

We discovered that 421 bots had been included in the Demographics Dataset because their behavior did not, at the time of the original collection, meet the criteria we had set for exclusion. They had been acti-

vated after the suspensions began and were tweeting at a high volume in January. Because of this—and the complexity of mid-stream name changes and suspensions—there also was a 0.02 percent difference in the number of accounts in the two comparison sets. The total number of accounts compared between the periods was approximately 18,500.

During collection of comparison data in January 2015, 72 accounts either changed their names or were suspended. Approximately another 30 disappeared in between the time we identified accounts that had not been suspended and our collection of tweet performance data, a relatively short period. With these substantial caveats in place, we made the following comparisons:

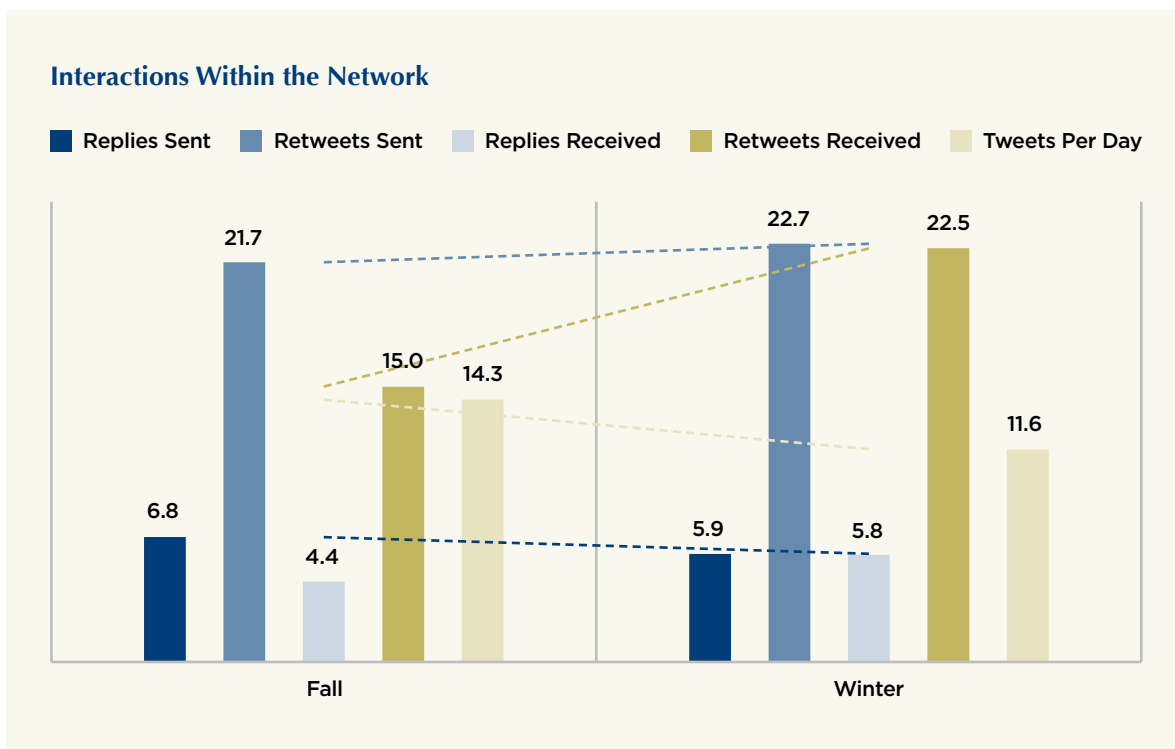


FIGURE 17: Change in performance, non-suspended ISIS supporting accounts, Fall 2014 versus January 2015

The average user tweeted less in January, despite heightened activity around the Hebdo attack during the collection cycle. The interactions shown in Figure 17 are delimited to only those taking place within the set (i.e., one user in the set retweeting another user).

As would be expected when key influencers are continually being removed from the ecosystem, users showed signs of being more internally focused on network members, with more retweets of content produced by other members of the network. These changes were relatively modest, however, likely because suspensions were already underway when collection started.

As expected, the average number of friends and followers per account increased modestly. Twitter ac-

counts which remain active almost always add followers, barring manipulative activity. The number of protected accounts in the network rose from 407 to 495, likely due to suspensions.

These metrics are of limited utility in understanding changes to the overall performance, because of the lack of a directly comparable set. Using an iterative process of collection would make it possible over time to create more valid comparisons.

In light of these limitations, we also examined the performance of tweets that originated with members of the network. In-network retweets dropped by almost 19 percent, while out-of-network retweets increased by just over 3 percent.

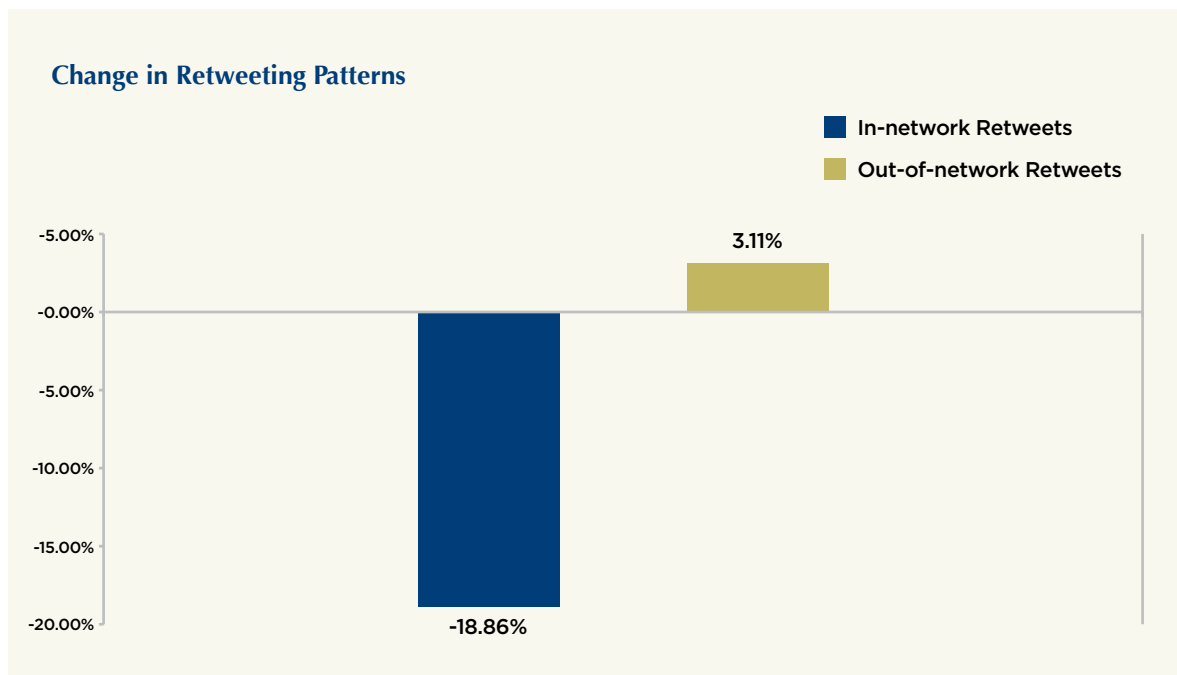


FIGURE 18: Change in performance of tweets by ISIS supporters, Fall 2014 versus January 2015

2.5.4 Partial comparison data

While we could not recreate the entire collection process in the time available, we produced a limited comparison set using the iterative process described above. It became apparent immediately that the composition of the network had changed.

The original 454 seed accounts (described in section 3.1) had followed almost 47,000 accounts (Level 1). Despite repeatedly broadening the criteria for seed accounts, we were unable to create a new Level 1 set of the same size in the time available.

This strongly suggested that the network had been substantially reduced and that the updated network

was much more internally focused, meaning ISIS supporters were increasingly following other ISIS supporters rather than a broader selection of accounts.

This could also suggest that it was more difficult for ISIS supporters to discover other ISIS supporters to follow. Twitter's recommendations engine influenced this dynamic. In 2013, a study by J.M. Berger found that Twitter's "who to follow" recommendations made it easy to discover jihadist accounts, once a user started following even one or two jihadist supporters.¹⁷

In February 2015, we recreated that experiment using ISIS-supporting accounts and found that the "who to follow" recommendations pointed to virtually no other jihadist accounts. However, Twitter later sent recommendations via email that accurately suggested ISIS supporters to follow. While the recommendation process took longer than it did in 2013, the net result was very similar.

We seeded a new collection using 550 accounts from the original set that were still active in January and February 2015, and after weeding for bots and accounts with more than 50,000 followers (the same criteria used for the Demographics Dataset), just under 25,000 accounts remained at Level 1.

All figures described in this paper as in-network metrics (referring primarily to interactions or in-network friends and followers) pertain to the overall dataset collected. For example, to calculate the number of replies received within the network of the fall collection, we defined "within the network" as the entire Census Dataset of about 47,000 accounts. However, the *average* number of replies received is the average of this full-set figure among only the accounts in the Demographics Dataset.

For the February dataset, similarly, the in-network calculations apply to the 27,000 total accounts we collected, prior to weeding. The averages for February are calculated using the full-set figures but averaging the totals only for the weeded set.

Due to time constraints, we did not collect information about which accounts the Level 1 users followed, and thus could not recreate the sorting metric (section 3.6). A casual examination of the set of 23,000 showed a very high incidence of ISIS supporters throughout the set.

Quantifiably, 660 accounts out of the 25,000 self-identified as ISIS supporters using the quick-code criteria, comparing very strongly to the 602 self-identified accounts in the Demographics Dataset of 20,000.

These data points strongly indicate that the ISIS-supporting social network on Twitter has been significantly constrained by the suspension campaign. We believe it is likely that the current size of the ISIS-supporting social network is lower than the fall 2014 estimate, and the estimated minimum could be as low as 30,000 at the time of writing.

The comparison set also revealed evidence that ISIS supporters were wearying of the battle with Twitter, using the number of new accounts as a barometer. While the spree of suspensions in September 2014 correlated to the creation of nearly 3,500 new accounts (section 1.7), the comparison data showed a steep decline in new account creation in the following months. This suggests that ISIS supporters were less often creating new accounts to replace those suspended, although it may also pertain to varying levels of suspension activity, which we were unable to estimate.

Both the original and the February 2015 analyses found similar numbers of accounts that had been created in September (accounting for the suspension trends described above), further reinforcing our belief that this more limited dataset is fairly representative of the overall current state of ISIS Twitter.

17. J.M. Berger, "Zero Degrees of al Qaeda," *Foreign Policy*, 14 August 2013, <http://foreignpolicy.com/2013/08/14/zero-degrees-of-al-qaeda/>.

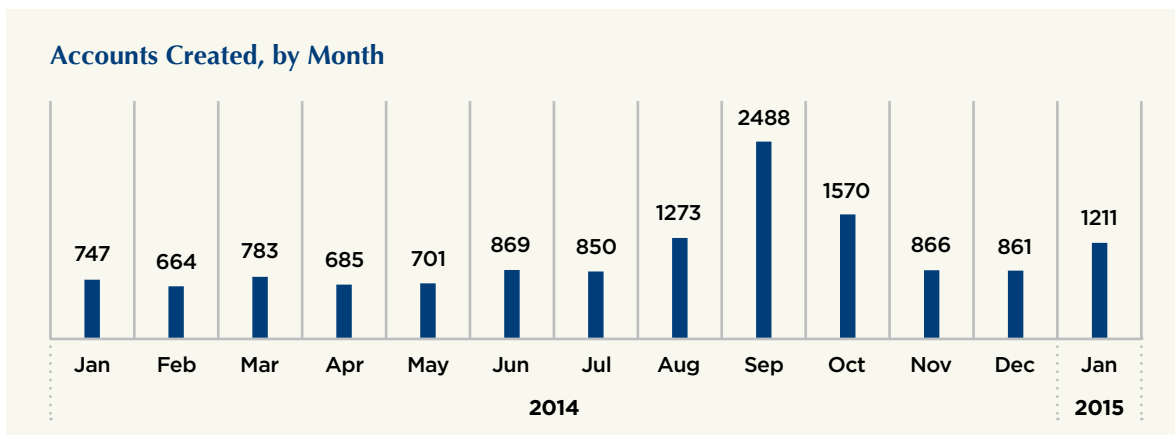


FIGURE 19: Month of account creation, February 2015 collection sample

Retweets received from other members of the network averaged 11.3 per user, compared to about 19.4 per user in the Demographics Dataset, a significant decline. However, we again stress that the two sets are not perfectly comparable. Received replies in the February set averaged 3.2 per user, compared to 5.19.

Tweets per day averaged 18.8, compared to 15.5 in the Demographics Dataset. This is significant in light of the decreased interactions described above. Users in the February set tweeted more, even as they recorded fewer interactions. However, the additional noise in the February set may distort these figures somewhat.

In the wake of Twitter’s suspension campaign, users responded by aggressively promoting new accounts that had replaced suspended accounts. Typically, tweets promoting ISIS accounts included a list of accounts with descriptive text. Our analysis isolated these as “non-reply mentions”—the inclusion of a Twitter handle in a tweet that is not a reply to another tweet.

In the February collection, users sent an average of 17.4 non-reply mentions to other users in the network, but received only 12 in return. In the Demographics Dataset from fall 2014, users sent an average of 12 non-reply mentions and received 10.4.

To produce a cleaner comparison, we also looked at these figures for accounts with fewer than 20,000 followers in both sets, expected to result in a higher proportion of ISIS-supporting accounts per the earlier analysis. The noise in the February dataset was especially concentrated among high-follower

accounts. This figure was particularly telling. In the February collection, these users sent an average of 40.3 non-reply mentions to other users in the network, but received only 8.6. In the Fall 2014 collection, users sent 11.9 non-reply mentions and received 9.8.

This data yields two major implications:

1. The amount of activity devoted to rebuilding the Twitter network (as opposed to disseminating propaganda, recruiting and other activities) more than tripled.
2. A five-to-one disparity opened up between the average non-reply mentions sent and the average received, indicating that the greatly increased activity was directed at fewer accounts. In other words, users were working much, much harder to promote the top accounts in the network. The increased amount of promotional activity did not correlate to an increase in the number of in-network accounts being promoted.

An interesting result running counter to the overall trend of depressed activity in the network pertained to follower counts. For the sets delimited to fewer than 20,000 followers each, the average number of followers nearly doubled from fall 2014 to February 2015—from 725 to 1408—with significantly fewer accounts recorded as having less than 200 followers, when compared to the Demographics Dataset. A possible explanation is discussed in section 4.3, which examines some of the potential tradeoffs incurred when degrading the network’s overall performance.

3. Methodology

Many claims regarding ISIS social media are based on data and methods that are—at best—unclear and often completely undisclosed. Therefore, we decided at the outset of this project to disclose how we collected and analyzed data in substantial detail, so that readers can understand how the findings were derived and to inform future research.

3.1 Starting Point

Co-author J.M. Berger spent much of 2014 identifying and collecting Twitter accounts for ISIS supporters using a variety of methods, including social network analysis and manual selection.

This starting list consisted of approximately 4,700 accounts around the time that the project began. Dozens were identified as particularly important supporters and members of ISIS's dedicated media and social media teams.

In the summer of 2014, ISIS's official accounts were suspended by Twitter. Several proxies for the official accounts were added to the list over time and monitored.

This starting dataset was monitored by reading tweets by users on a near daily basis. This method was revised over the course of several months, during which some users were weeded from the set after a review of their content suggested they were not unambiguous ISIS supporters.

Examples of non-supporters who show up in ISIS social networks include a relatively small number of journalists and academic researchers; accounts whose interests overlap with ISIS in an adversarial way, such as accounts controlled by supporters of other jihadist factions in Syria, including Jabhat al Nusra and Ahrar al Sham; and online jihadist activists associated with al Qaeda and its affiliates.

At the outset of the project, the most engaged accounts in the dataset were identified using a refined version of the methodology outlined in Berger and Strathearn's study "Who Matters Online: Measuring influence, evaluating content and countering violent extremism in online social networks."¹⁸

Data was collected by proprietary software originally coded by Dan Sturtevant using guidelines set

out by J.M. Berger, with collection and analytical revisions specific to this project written by Jonathon Morgan.

The initial process of identifying a larger set of ISIS supporters involved collecting all of the accounts followed by the seed accounts ("friends" of the seeds).

Our guiding analytical principle was that an ISIS supporter online could be best defined as someone who was followed by at least one other ISIS supporter, rather than someone who tweeted specific kinds of content within a particular time frame.

This approach was taken in part because of an observed trend of ISIS supporters deleting or protecting their tweets, but also because extremists—foreign fighters especially—may be subject to external pressures that cause them to go silent, such as lack of Internet access or intelligence agency scrutiny. Some also use Twitter purely for private messaging or covert signaling, a behavior that has been observed. Barring a few truly rare covert accounts, an ISIS supporter of any significance in the contexts above should at some point be followed by at least one other ISIS supporter.

In contrast, follower lists tend to be filled with observers, including adversaries, spies, and researchers. Passive accounts used by researchers and analysts, for instance, might follow many ISIS supporters and be followed by few or none. By limiting the direction of collection to following lists, we hoped to minimize the impact of such noise.

Following general observations of Twitter activity, we decided to limit the seed accounts based on the number of accounts each candidate followed. Users employ different strategies to following other users on Twitter. Some follow everyone who follows them, either as a courtesy or as a social marketing tactic. Others follow a wide variety of accounts in order to obtain information from many sources;

18. J. M. Berger and Bill Strathearn, "Who Matters Online: Measuring influence, evaluating content and countering violent extremism in online social networks," *International Centre for the Study of Radicalisation and Political Violence*, March 2013, http://icsr.info/wp-content/uploads/2013/03/ICSR_Berger-and-Strathearn.pdf.

however, extremists tend to rely more on news from within their own circles.

Based on anecdotal observation and the amount of time required for collection (see section 3.2), we designated accounts following 500 or fewer others as optimally focused and least likely to produce noise in the form of non-ISIS-supporting accounts.

We also limited the seeds to accounts that had tweeted within four months preceding the start of collection in September. We did this to eliminate a very large number of accounts that had tweeted only through a known ISIS Twitter strategy, the use of the “Dawn of Glad Tidings” auto-tweeting app (described in section 1.13).

In the end, a final analysis of the candidate accounts was performed in late September and early October 2014, and a total of 454 accounts were selected to serve as “seeds” for the ISIS Twitter Census dataset. We collected all of the accounts they followed, for an initial Census Dataset of 49,379 accounts.

A quick review of the data indicated that ISIS supporters were extremely unlikely to have more than 50,000 followers, likely the result of Twitter undertaking a substantial initiative to suspend ISIS supporters starting in September. An initial review of accounts with more than 50,000 followers revealed that they were obviously and overwhelmingly related to topics other than ISIS.

In light of this, all accounts with more than 50,000 followers were removed from the Census dataset after initial data and some simple metrics were calculated.

3.2 Challenges and Caveats

A number of unusual factors complicate analysis of ISIS Twitter networks. Although some factors are related to operational security, Twitter’s campaign of suspending ISIS supporter accounts was respon-

sible for the most significant challenges, either directly or indirectly. A brief list of some of these complications should highlight important caveats on this research.

- **Suspensions:** The number of ISIS accounts being suspended by Twitter on a regular basis appears to be much higher than previously reported. Because our data collection and analysis methods require time to complete, accounts were suspended and created in the middle of the collection process, skewing results.
- **Name changes and voluntary deletions:** In an effort to avoid suspension, some ISIS users changed their screen names or deleted their accounts during the process of analysis, skewing results. Our analysis was generally able to account for name changes, but some distortions may have occurred.
- **Private accounts:** For reasons of operational security or as a hedge against suspensions, some users either temporarily or permanently set their privacy settings to “Protect Tweets,” which prevented analysis of their content and inhibited, but did not entirely prevent, analysis of their relationships. While this tactic was much discussed in circle circles of ISIS users, it was not observed to be widely deployed in practice.
- **Multiple accounts by the same user:** In some cases, ISIS supporters who anticipated being suspended maintained multiple backup accounts. We did not attempt to eliminate these accounts from our analysis. Based on anecdotal observation, they likely represent less than 1 percent of the total collected; additional investigation might provide a more definitive estimate.
- **Non-organic noise:** Some users employed guidelines published by ISIS social media strategists, instructing them to follow large numbers of accounts which were not fully relevant to ISIS. The purpose of adding this noise to the system was generally to avoid suspension, but it may have a related function of complicating analyses such as this one.

- **Tweet deletion:** A relatively small number of ISIS-supporting Twitter users employ services that allow them to delete all of their tweets, or were noted to delete tweets manually.
- **Deceptive practices:** ISIS uses several practices designed to amplify its apparent support on Twitter, including “bots” (computer software that creates activity on a social media account in the absence of a human user) and spam (purchased tweets promoting ISIS content). We were able to compensate for this, although not perfectly.
- **Size of the dataset:** Certain practical considerations come into play when dealing with a very large dataset, primarily restrictions on the speed of data collection speed. Twitter enforces such restrictions regarding friend and follower relationships. Some of these issues are detailed in Appendix B. Over the course of data collection, the dataset changed composition between the start and end of the process, complicating analysis. For instance, in addition to the suspensions noted above, most accounts added followers during collection.

Each of these limitations can be compensated for using analytical techniques. Some are more difficult to offset than others, and each one adds a certain amount of noise to the final dataset even after compensation. Some of these effects and compensating analytical approaches will be described more fully below.

3.3 Bot and Spam Detection

ISIS is known to use a variety of deceptive techniques to inflate the appearance of support online, including apps, bots and spam services (see comments and definitions above under ISIS Supporter Basic Metrics).

At one point, ISIS employed a single, widely adopted app that automatically tweeted information out from thousands of ISIS-supporter accounts.¹⁹

For some of these accounts, the user also tweeted normally. Others only sent tweets generated by the app. All of the latter accounts were simple to remove from our analysis.

Since Twitter terminated this app in June 2014, ISIS supporters have implemented new apps. Unlike the original app, these use a variety of techniques and are not subject to a single point of failure that would enable them to be silenced en masse. The tradeoff is that the new apps and bots are also less synchronized and therefore less powerful (see section 1.13)

The metrics we applied to rank accounts (section 3.6) naturally downplay such deceptive techniques, but do not totally eliminate them. Therefore we took steps to eliminate accounts that were primarily controlled by bots and apps from the set.

A variety of techniques used in past research by J.M. Berger have proven effective in identifying bots and were applied to this process. The two most important in this context were identical tweets and Twitter platform information.

Identical tweets are tweets that contain identical text but are not retweets. A large number of an account’s tweets that are identical to other tweets in the dataset is a strong indicator of a bot or spam account. Through manual spot-checking of accounts and correlation with Twitter platform information, we estimated that accounts with more than 30 identical tweets among the tweets collected were most likely to be bot or spam accounts, and were thus removed from the dataset.

Twitter platform information is associated with each individual tweet, indicating whether the tweet was sent from the Twitter web page or from a Twitter client such as Tweetdeck. Because this can be falsified, platform information alone was not an adequate indicator. Many tweets were sent using apps that clearly identified themselves as bots, spammers, or automated services. Accounts that sent most or all of their tweets using such platforms were removed from the dataset.

19. Berger, “How ISIS Games Twitter.”

These conservative steps did not completely eliminate bots and spammers from the data. 421 more accounts were discovered to be bots in a second collection of tweets in January 2015; they were not detected in the first round of eliminations because they had not been fully activated.

Slightly more than 3,000 accounts were eliminated from the Census Dataset as likely bots or apps prior to the scoring and selection of the Demographics Dataset. This conservative approach allowed for the retention of some users who employed apps in addition to tweeting normally. Many of the remaining bots and spammers were ranked as less likely to be ISIS supporters by the natural function of the other metrics.

3.4 Description of Data Collected

Armed with the 454 seeds, we proceeded to collect user information and the 200 most recent tweets of all accounts followed by the seed accounts, as well as user IDs for accounts followed by the new group of users. We distinguish between the different types of information collected in the following ways:

- The seed accounts were designated Level 0.
- The people they followed were designated Level 1.
- The people followed by Level 1 users were designated Level 2.

- Accounts mentioned by Level 1 users who were not part of Level 1 were designated as Level 3. Some Level 3 members were later determined to be members of Level 2 as well.

For Levels 0 and 1, all user information, tweets, and following/followed relationships were collected. For Level 2, only numerical user IDs were collected at the outset, but the data collected at Level 1 indicated how many Level 1 accounts followed each Level 2 account. For Level 3, only Twitter handles were collected.

There were almost 2 million users at Level 2, presenting an obstacle to collecting data in a timely manner. In December 2014, we collected user information for all Level 2 users that were still active. We did not collect following/followed relationships or tweets, but we were able to infer some information about those data points from the data previously collected.

After the initial rounds of data collection described above, we analyzed the results and experimented with several different approaches to sorting and classifying the accounts for which we had collected adequate data.

The tables below outline the data collection and level classifications.

Network Level	Definition of Level	Data Collected
<i>Level 0</i>	Initial seed accounts selected	All user and profile information, last 200 tweets of each user, list of users followed by each seed account
<i>Level 1</i>	All accounts followed by Level 0 accounts	First round collection: All user and profile information, last 200 tweets of each user Second round: All user and profile information, last 200 tweets of each user, list of users followed by each Level 1 account
<i>Level 2</i>	All accounts followed by Level 1 accounts that were not already found in Level 0 or Level 1	Second round: Numerical user IDs collected Third round: User information collected
<i>Level 3</i>	Accounts mentioned by Level 1 users who were not part of Level 1	First round: Only Twitter handles collected Third round: Some user information collected

TABLE 1: Classifications of data collected

Data Collected	Definition	Amount
<i>User information</i>	Includes username, display name, numerical ID, followers, following (friends), URL of profile image, total number of tweets over the lifetime of the account, date that the account was created, profile text, location (if provided), web site link (as shortened by Twitter) and other data	This data was collected for more than 2 million accounts
<i>Tweets</i>	Last 200 tweets were collected for Level 0 and Level 1 users. Includes text of tweet, time of tweet, location information (if provided), platform used to send tweet	This data was collected for just over 49,000 accounts
<i>Following relationships</i>	The accounts followed by a Twitter user	This data was collected for just over 49,000 accounts

TABLE 2: Description of data collected

3.5 Data Codebook

Samples of Twitter accounts from the collected data were coded to evaluate which ones were ISIS supporters and which were not. In cases where an account's standing was ambiguous, it was weighted as half of a full supporter. The following criteria were applied, in the following order, to determine whether a user was an ISIS supporter:

- The content of the user's tweets were examined to see if they were unambiguously for or against ISIS. If tweet content included both ISIS and non-ISIS jihadist content, or was strongly inconsistent with ISIS's stated values and ideology, the user was coded as a non-supporter.
- If the first step did not produce a conclusive result, a user's profile and profile images were examined for indications the user was an ISIS supporter (i.e., using the ISIS black flag or displaying pictures of ISIS emir Abu Bakr al Baghdadi).
- In the absence of both tweets and profile identifiers indicating a clear preference, the user's list of followers was examined. If, on subjective examination, a significant proportion of their followers used profile markers to clearly indicate support for ISIS in the absence of anti-ISIS users (such as members of Jabhat al Nusra), the user was coded as an ISIS supporter.
- If all other criteria failed, we examined who the user followed. If that list was substantially

focused on ISIS supporters in the absence of a meaningful number of non-supporters, it was coded as a supporter.

- If an account was suspended, it was presumed to be an ISIS supporter, given the context of the collection process. At the time of collection, Twitter suspensions were highly focused on ISIS. If we were repeating collection in the current environment at the time of writing, we would likely consider a different approach, since Twitter has begun suspending non-ISIS jihadists more aggressively in the interim.
- If an account was deleted by the user and not by Twitter, it was scored as a 50 percent probability the user was an ISIS supporter. However, there is reason to think that total may be higher as many ISIS users deleted their accounts around the time Twitter began its crackdown.
- We also performed a quick-coding process in which all users in the Census dataset who explicitly identified with ISIS in their user handles or profile information were identified by keyword searches. Keywords included ISIS, _IS or IS_, the ISIS slogan *baqiyah* and variations thereof, as well as references to the Islamic State or the caliphate using common phrases in both English and Arabic. These results were vetted for obviously non-responsive accounts (such as strings using "thisis" [this is] or accounts whose handles indicated opposition to ISIS. The quick-coding results were primarily used to measure the overall effectiveness of the metrics (section 3.6) rather than the accuracy of the overall set.

We examined and coded a random sample of 1,000 Twitter accounts extracted from a slightly smaller dataset derived in a method identical to the main Census dataset (and with substantial overlap). Within the random sample, 65 percent of accounts were clearly ISIS supporters, 5 percent were ambiguous, and 30 percent were non-supporters. We relied on this number as a kind of control group, in the sense that results which appeared dramatically out of line would be scrutinized carefully for methodological errors.

Coded results from the random sample were integrated into the main Census Database. We additionally coded ranges of accounts according to ranking using the metrics in section 3.6, in order to gauge the effectiveness of the metrics at different levels.

For accounts coded as ambiguous, an additional review was made of all metrics collected to see if the account in question could be further resolved. In the final codebook, 4 percent of all coded samples were recorded as ambiguous.

3.6 Sorting Metrics



FIGURE 20: Responsive accounts (blue bars) ranked left to right by engagement method as detailed in “Who Matters Online”




FIGURE 21: Ranked by ratio of in-network-followers to out-of-network followers



FIGURE 22: Ranked by metric combining weighted clique membership, in- and out-of-network followers, and engagement

After eliminating all likely bots and all accounts with more than 50,000 followers from the Census dataset, we were left with 43,538 accounts to evaluate as ISIS supporters or non-supporters.²⁰ We experimented with several different methods of sorting the Census Dataset according to the probability that an account belonged to an ISIS supporter.

We evaluated each of the metrics using visualizations of accounts quick-coded as ISIS supporters within the refined Census dataset of 43,538 accounts. In Figures 20, 21 and 22, responsive ac-

counts are represented by blue bars, and non-responsive accounts are represented by white bars.

Each chart is ranked left to right, meaning that bars to the left scored higher using the designated metric than bars to the right. A concentration of blue bars on the left of the chart indicated the metric was more successful at sorting relevant accounts than one in which the blue bars were spread out over the whole chart. The blue bars are slightly weighted for more convenient viewing, slightly overrepresenting the proportion of responsive accounts in visual terms relative to numerical precision.

20. We used a very similar, slightly smaller set for coding purposes and to evaluate the effectiveness of metrics, which performed identically among the two sets. The final demographics were derived from the more complete set. The differences between the two sets related primarily to how bots were eliminated, with the final set representing a more conservative estimate of the number of likely bots that could be safely excluded.

We experimented with more than a dozen sorting metrics. Three are depicted in the charts above. The first metric was the engagement metric described at greater length in the “Who Matters Online” paper. As can be seen in Figure 20, this produced a better than random result, but responsive accounts were still significantly distributed throughout the set. The engagement metric tends to produce a very relevant list of accounts at the highest-scoring levels, but many relevant accounts are still found at lower-scoring levels.

At the outset of the project, we intended to examine two additional properties for their merits in sorting relevant accounts: **cliqueishness** and **in-network focus**. Both showed significant strengths and weaknesses individually; when combined they proved more powerful than any of the other approaches we tried.

Cliques are network structures in which every node is connected to every other node. In a social network, this is most easily understood as a group of users in which every member of the group follows every other member.²¹ As we define it, cliqueishness represents a Twitter account’s proximity to cliques in the network, measured by the number of a user’s friends who belong to cliques.

When a network reaches even a modest size, the number of cliques it contains quickly becomes too large to compute.²² Therefore we had to impose limits on our calculations. We tried different approaches to generating cliques, including an iterative-deepening approach (building out from the seed accounts) and a top-vertices approach (building out from the 100 most-connected entities in the graph). As neither of these proved entirely satisfactory, we opted instead to use what we defined as in-network focus to identify the most useful starting points to search for cliques.

In-network focus indicates whether a member of the network follows more accounts within the Level 0 and Level 1 network than outside (total number of friends minus in-network friends). Since Level 0 was

comprised entirely of ISIS supporters and Level 1 was heavily slanted toward ISIS supporters, we theorized that users whose networks were more concentrated internally would be more likely to be ISIS supporters.

We also chose the 100 most internally focused accounts as a starting point to generate cliques.

As seen in Figure 21, the ratio of a user’s in-network friends to out-network friends resulted in some very accurate hits at the top, but with many responsive accounts near the bottom as well. These represented users whose network connections were primarily outside the Level 0/1 network but who were still ISIS supporters. Overall, while the results had some merit, an unacceptable number of ISIS supporters scored too low using only the in/out ratio to serve as a primary metric.

After a considerable amount of trial-and-error experimentation, we developed the metric visualized in Figure 22, the most effective approach we tried. While some other metrics were equally or nearly as effective within the top 20,000 accounts, the metric shown in Figure 22 was more effective at lower ranges as well. The formula combined and weighted many of the previously successful metrics and thereby produced the strongest results.

The approach added the total number of each user’s:

- In-network friends;
- In-network followers multiplied by eight for weighting purposes;
- The sum of the in-network friends of all the users in the largest clique to which the user belonged, multiplied by 16;
- We added one metric that may be peculiar to this dataset. In the wake of Twitter’s suspension campaign, users responded by aggressively promoting new accounts. Typically, tweets promoting ISIS accounts included a list of accounts with descriptive text. Our analysis broke these out as “non-reply mentions”— the inclusion of a Twitter han-

21. Robert Hanneman, “Introduction to social network methods, 11. Cliques and sub-groups,” University of California, Riverside, http://faculty.ucr.edu/~hanneman/nettext/C11_Cliques.html#subgraph.

22. “The Maximal Clique Problem,” DNA Computing, <http://cs.stanford.edu/people/eroberts/courses/soco/projects/2003-04/dna-computing/clique.htm>.

dle in a tweet that is not a reply to another tweet. We weighted these mentions by multiplying the number of non-reply mentions that an account received by four. Because ordinary Twitter users also employ non-reply mentions—for example, when recommending accounts to follow—so this metric may also apply to non-ISIS networks, but would likely require a different level of weighting (see Control Group, Appendix B).

The sum of these measurements was then multiplied by the ratio of the user’s followers within the network (Level 0/1) divided by the number outside the network (total number of followers minus in-network followers).

We dubbed the metric IQI (for its major components, in-network, clique, and interaction).

3.7 Metrics Performance and Estimates

The IQI metric performed accurately at several ranges, allowing an estimate of the total number of ISIS supporters in the Census dataset. The table below shows the breakdown of accounts as measured by a random sample.

The sample was slightly clumpy, resulting in uneven distribution, with a disproportionate number of accounts randomly falling into 1 to 10,000 range and the 30,000 to 40,000 range, biasing the results for the top 20,000 and the lower 23,000 slightly higher. Part of the reason for this incongruity is the difference in size of the dataset used for random coding, which was collected earlier than the main census dataset and therefore had fewer members.

Ranking Using IQI Metric	ISIS Supporters	Combined Coded Results in Range
1 to 200	98.5 percent	
1 to 5,000	94.3 percent	
1 to 10,000	94.8 percent	94.5 percent
10,000 to 20,000	93.4 percent	
20,000 to 30,000	77.1 percent	
30,000 to end	29 percent	50.9 percent
TOTAL	~31,012 responsive accounts (71.2 percent of total set)	

TABLE 3: Accuracy of metrics based on random sample of 1,000 coded accounts taken from entire set

In order to produce an evenly distributed sample and understand the trend produced by the metrics, we also coded accounts in consecutive panels of 200 at different ranges within the set. The panel from 20,000 to 20,199 was included calculations for both the top 20,000 and lower 23,000 to even out the distribution of coded accounts.

The random sample was coded first; therefore accounts already coded were not recoded. While

coding guidelines were fairly robust, the researcher coding most of the ranged sample was slightly more conservative (i.e., less likely to designate any given account as a supporter) than the researcher coding the random sample, due to having more experience with the subject matter. When an account was designated as unclear, a second coder reviewed the account using additional metrics to see if it could be resolved. If this was not possible, the account was marked as 50 percent responsive.

Ranking Using IQI Metric	ISIS Supporters	Combined Coded Results in Range
1 to 200	98.5 percent	91.8 percent
5,000 to 5,199	93.7 percent	
10,000 to 10,199	88.3 percent	
15,000 to 15,199	86.8 percent	
20,000 to 20,199	76.8 percent	
30,000 to 30,199	48 percent	50.3 percent
40,000 to 40,199	26.3 percent	
TOTAL		74 percent

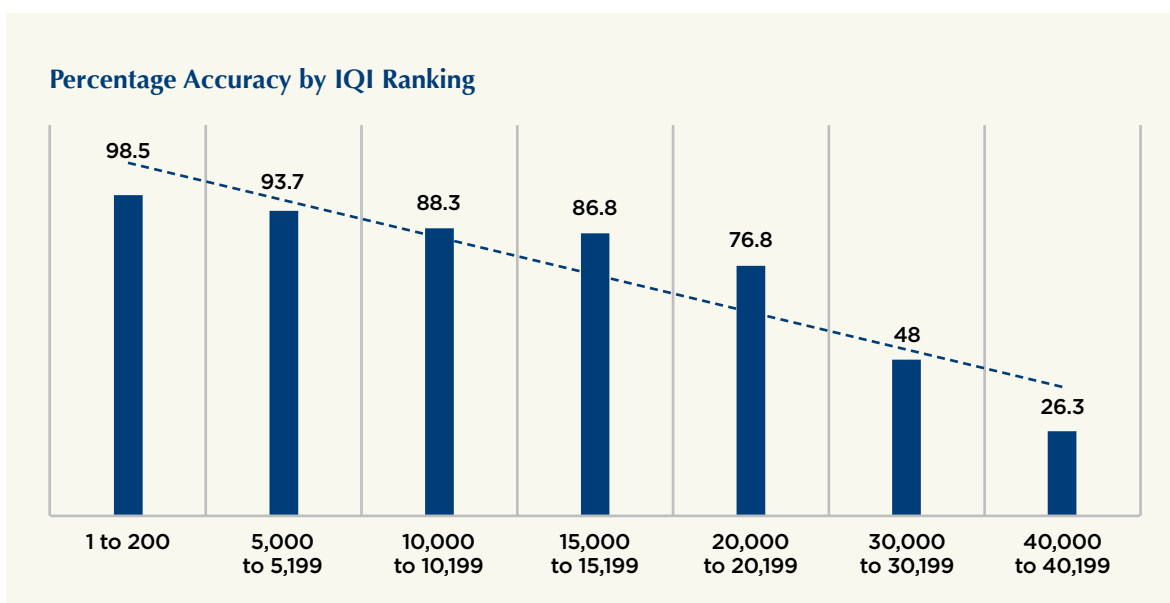


FIGURE 23: Accuracy of IQI metric according to range of scores (high to low)

While the ranged coded accounts provided superior distribution of the sample, the nature of the metrics—specifically a heavy weighting of cliques and non-reply mentions—may have produced streaks of unusually high or low values.

To test this theory, we coded a range of 300 accounts from a position slightly higher than the median of the Census Dataset (at and after the cutoff for inclusion in the Demographics Dataset, and where we might expect more ambiguous findings) and found evidence to support our theory (Figure 24). Within the run of 300, relevance increased as

the IQI metric decreased, despite the strong broader trend of positive correlation between relevance and IQI score.

A more detailed examination of the entire Census Dataset found that the IQI metric tended to produce runs of accounts with an identical score greater than zero in certain ranges after the top 2,000, as the input data became more uniform, which may also have contributed to the streaks. These runs typically lasted less than 10 accounts, but in rare cases were as high as 85. A long run of scores equaling zero was found in the lower

5,000 accounts, signaling that sorting in that range was essentially nonexistent.

Within the sample of 300 consecutive scored accounts, streaks of non-responsive or ambiguous accounts (such as accounts that did not overtly

support ISIS) appeared to correlate with higher numbers of either in-network friends or followers. This suggests such accounts represent sources of information or opinion that were highly relevant to ISIS supporters, even if they did not overtly support ISIS in tweets.

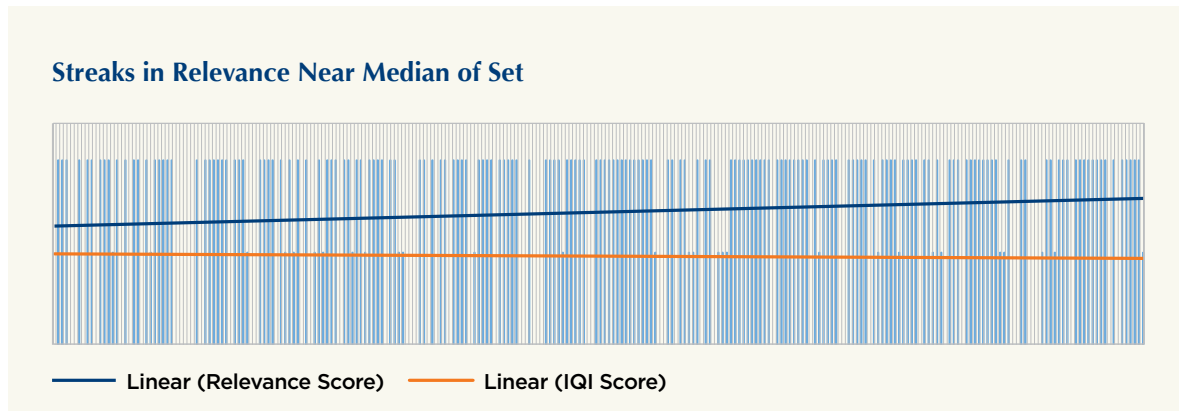


FIGURE 24: Streaks of ambiguous accounts were seen in the sample, due to the nature of the metrics. Blue bars represent ISIS supporters, white bars represent non-supporters, and half-bars represent accounts that could not be conclusively identified as supporters or non-supporters. The orange line represents the IQI score.

Merging all of the recorded values from both the random and ranged coded sets produced the following results:

Ranking Using IQI Metric	ISIS Supporters	Combined Coded Results in Range
1 to 9,999	96 percent	93.2 percent
10,000 to 20,000	89.9 percent	[Demographics Dataset]
20,001 to end	50.2 percent	50.2 percent
TOTAL		74.8 percent

TABLE 4: Accuracy of metrics, all manually coded samples

The final total for the overall set skews slightly higher due to additional ranged samples near the top. However, the metrics values flattened in the bottom half, somewhat offsetting the impact of the distribution imbalance. Our final estimate for the number of ISIS supporters in the Census Dataset is closer to 70 percent, in line with our initial “reality check” random sample. We estimate a rounded total of approximately 35,000 ISIS supporters within the Census Dataset.

Based on the figures from each analysis and the combined results, along with other factors noted above, we estimated that the top 20,000 accounts would consist of approximately 93 percent overt ISIS supporters, with a margin of error of about 2.54 percent. There was a strong potential that highly relevant accounts in the set did not meet our conservative criteria for designating a user as a supporter. We used the set of 20,000 set to analyze the demographic information presented earlier in this report.

3.8 Machine Learning Approach to Level 2 Results

In order to estimate the total number of ISIS-supporting accounts on Twitter, we had to deal with the 2 million accounts at Level 2, the accounts that were followed by Level 1 users but were not included in Level 0 or 1.

For Level 2, only user information could be directly collected, although data collected at Level 1 enabled us to see how many Level 1 users followed a Level 2 account, as well as whether they had referred to a Level 2 user in a tweet.

The data was insufficient to develop a complete picture of the number of ISIS supporters found in Level 2, which represents all users two hops away from the seeds at the time of the collection (friends of friends). Despite significant caveats for incomplete data, we tried several approaches and found wide variations in the results.

Twitter allows users to include a short personal description in their account profile, which is

often helpful in deciding whether or not a given user is an ISIS supporter. While it is not practical to read each of the over 2 million profile descriptions collected for the census, we taught a machine learning model to analyze their contents.

Machine learning models need to be trained on sample data that has already been categorized. We used a set of roughly 6,000 accounts that had been hand-coded as either ISIS supporters or non-supporters. By focusing only on users who had explicitly declared their support for ISIS in the profile description, we trained the model to recognize words and sentence fragments that strongly correlate with ISIS support.

The model correctly deduced that Arabic words like *succession*, *linger*, *Islamic State*, *Caliphate State* or *In Iraq* are often present in the profiles of ISIS supporters (Figure 25). While this may seem obvious to human readers, this is an important step in teaching a machine to read and categorize text. Note that the English translations in the chart are provided for reference, but the model was trained with Arabic text.

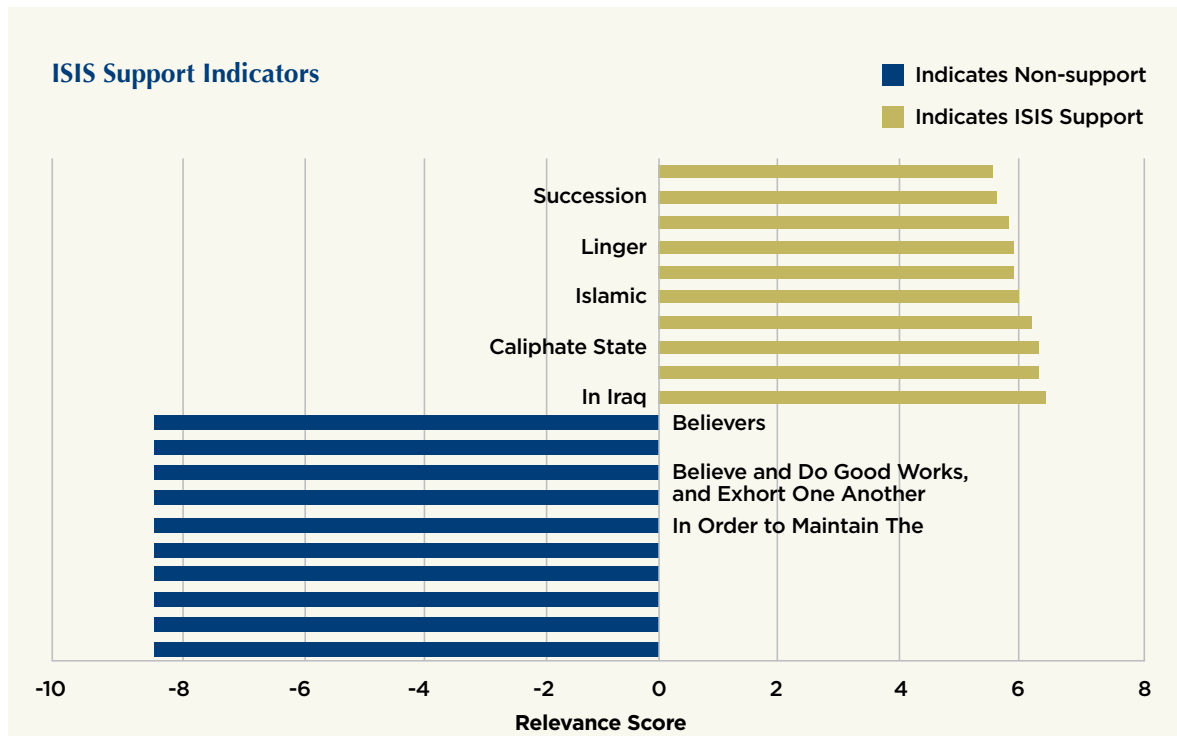


FIGURE 25: Selected phrases in an account profile that correlated to support or non-support of ISIS

The model was 94% accurate in distinguishing between the profiles of supporters and non-supporters when tested against a set of 1,574 hand-coded accounts. However, 25–35% of users collected in the census did not supply a profile description, so the overall likelihood of identifying a supporter or non-supporter using only profile classification is closer to 70%. Furthermore, the machine learning approach, when employed without other network-based constraints, produced a visibly noisy set that included many obviously non-responsive accounts.

3.8.1 Predicting Supporters

While useful, the machine learning model was not a silver bullet. Words can be misleading even when they are available. It is equally important to investigate the extent to which a user participates in networks of known or likely ISIS supporters. This concept of network participation is at the core of the IQI metric, which is the most accurate way we have found to identify ISIS supporters.

Unfortunately, due to data access limits imposed by Twitter, it would take months or even years to collect the values required to calculate that metric with a large dataset.

Instead, we were forced to estimate using a subset of the data, including but not limited to the profile classification score discussed above. We used another machine learning technique called feature selection to algorithmically determine which of the available metrics best correlate to a user being identified as an ISIS supporter or not.

The predictive model built using a wider range of features was only 80% accurate when tested on a set of 5,895 hand-coded accounts. For comparison, the IQI metric is 98% accurate on the same test set. Using this model, we would estimate between 65,000 and 95,000 ISIS supporters amongst Level 2 users.

However, relative to the criteria used in section 3.8 for estimating Level 2, this figure seems considerably higher than expected, as well as higher than anecdotal observations would support. A casual examination of some accounts that were scored as responsive found a relatively low proportion of overt ISIS supporters per our criteria.

Any discrepancy may be due in part to extrapolating the results of the machine model against the coded samples—which are taken from a set consisting of 65 to 75 percent ISIS supporters—to a group of 2 million, which even in the most aggressive view contains less than 5 percent. Additionally, many of the terms that correlated positively to ISIS support, such as caliphate and Islamic, are frequently used by non-supporters.

In light of these factors, we opted to sort the set of the 66,000 most-responsive accounts according to the ratio of their in-network followers (meaning followers in Level 0 and Level 1) to their total follower counts. This was evaluated using the same quick-coding method described in section 3.6, and yielded very encouraging results, if not as rigorously accurate as the full IQI method.

Based on the findings—which roughly conformed to several other less rigorous approaches to crafting an estimate—we project that there were a minimum of 16,000 ISIS-supporting accounts in the Level 2 dataset.

While the machine learning model left open the possibility that as many as 70,000 supporter accounts might be found at Level 2, we believe this figure is too high to be credible. We also cannot rule out a significantly higher total in Level 2 if the models we employed were too conservative; however, we have seen no observational evidence to support an estimate at even the highest range we have provided.

4.

Conclusions: Pros and Cons of Suspending ISIS Supporters on Twitter

This paper was motivated in part by a desire to better understand the impact of Twitter's suspension of ISIS supporter accounts on the performance and coherence of the overall network. An improved understanding would further inform the public debate over the value of suspensions. There are many stakeholders in this debate, including free speech advocates, counterterrorism officials, CVE programs, journalists, and independent researchers.

The most important stakeholder in this debate is Twitter, the for-profit company that owns the social media platform and likely issues most of the decisions about which accounts will be suspended. However, none of the stakeholders—not even Twitter, which appears to be subject to government requests for monitoring and mediating content—fully own the issue.

Each stakeholder has different interests in the problem. Counterterrorism professionals, for instance, may be highly interested in open-source intelligence about ISIS, whereas Twitter is not in the business of counterterrorism and has an interest in protecting the privacy of its users. Some CVE programmers wish to project anti-extremist messaging into the ISIS space, while others prefer to see potential avenues of radicalization—including Twitter—closed down. Conversely, many open-source analysts following ISIS want easy access to information in order to avoid the elaborate and sometimes expensive procedures required to find less visible sources.

Three crucial questions surround the debate over suspension of terrorist social media accounts in general, and ISIS accounts in particular.

- Is it ethical to suppress political speech, even when such speech is repugnant?
- Do suspensions destroy valuable sources of intelligence?
- Do suspensions have a detrimental effect on targeted networks?

While the first cannot be directly addressed by this study, our research holds implications for the subsequent two questions. Debates over these issues have thus far relied heavily on anecdotal observations, strongly held opinions, and small data samples derived with relatively weak—or entirely undisclosed—methods. In these areas, this report can add contribute useful evidence to inform the debate.

4.1 Intelligence Value

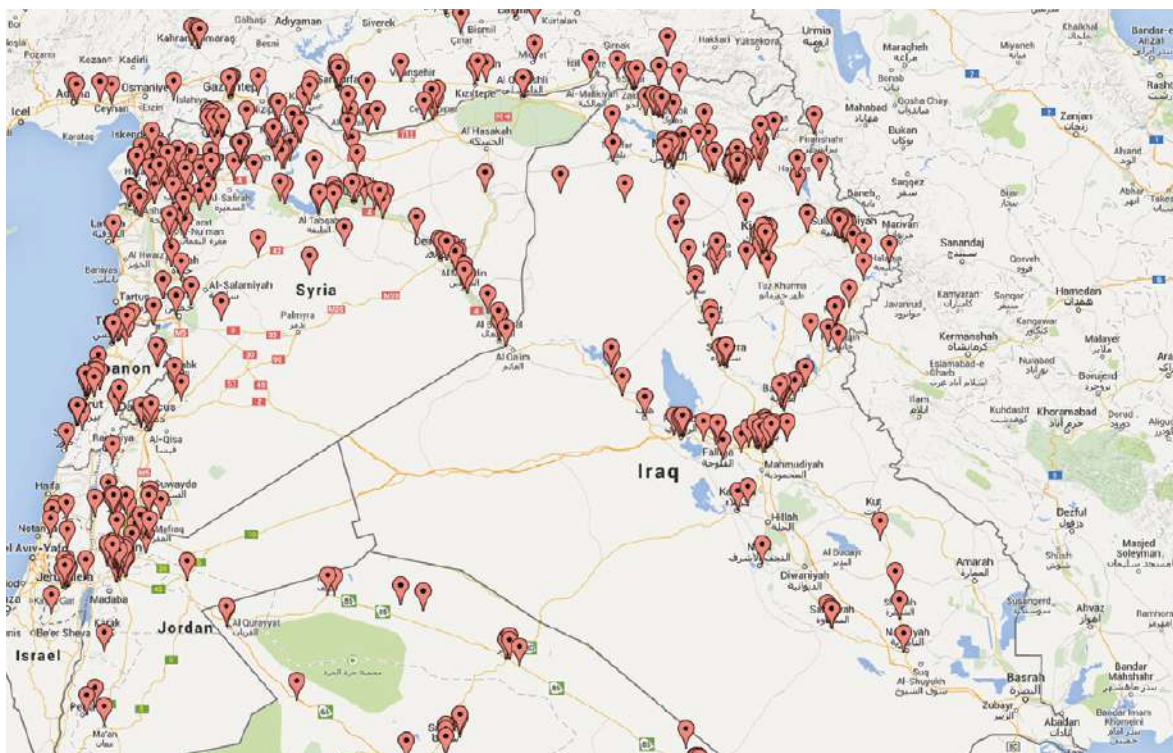


FIGURE 26: GPS-located tweets in Iraq and Syria, late December 2014

There is clear intelligence value to be extracted from the ISIS accounts we examined. Although the volume of material created challenges in approaching this material systematically, the data analysis provided a number of clear insights. Most prominently, a significant number of accounts provided reliable GPS coordinates in ISIS territories. A subsequent data collection in late December 2014 detected even more relevant GPS coordinate data in Iraq and Syria, even as ISIS was warning its members against the practice.

There are a number of intelligence applications for such data. For example, when we correlated the GPS data to other kinds of inferred geolocation, we were able to identify Twitter profile categories where users are likely to lie or provide misleading information. In another example, users who are GPS-located to Iraq and Syria can be subjected to further social network analysis to estimate which of their online friends are also located there. This is only the tip of the iceberg.

The volume of collected tweets and the number of languages employed by users created prohibitive obstacles to an algorithmic evaluation of the content posted within the time and budget of this project. We could, however, load the most relevant accounts in the Demographic Dataset into a Twitter client and monitor the users over time, continuing a process that has involved thousands of ISIS-supporting accounts over the course of 2014.

We have observed anecdotally that smaller accounts often post material regarding local events as they were happening, and that medium-sized accounts often provide an early glimpse at ISIS media releases.

Additionally, by monitoring the combined feeds on a large-screen TV, it was often possible to evaluate—literally at a glance—the importance of new media material, themes, and issues. When ISIS releases important propaganda (such as the video showing the killing of Jordanian pilot Muaz al-

Kasasbeh), the video can be seen propagating in a distinctive visual pattern on the screen, enabling a swift evaluation of its significance and whether it was an official release.

Additional monitoring—alongside network analysis to refine subgroupings in the set—would certainly improve this intelligence value. ISIS has taken steps to ensure its operational security on social media, but it cannot accomplish its propaganda, recruitment, and operational missions on Twitter without exposing itself to scrutiny.

Our anecdotal observation indicates that the most valuable intelligence tends to emanate from the least obvious vectors, such as accounts with very small numbers of followers. The most active and visible accounts contain more noise, and their content is more carefully stage-managed by ISIS and its adherents. Based on both anecdotal observation and ISIS social media strategy documents, original information tends to flow from more obscure accounts to more visible accounts such as the *mujtahidun*, the core group of ISIS supporters devoted to disseminating information that originates elsewhere.

The ability to accurately identify tens of thousands of ISIS supporters on Twitter provides ample room for the suspension of accounts that have strong operational, recruitment, or propaganda value, with little or no functional loss of intelligence.

From a purely instrumental counterterrorism point of view—without regard to the many other issues at play—the challenge is to sufficiently degrade the performance of the network to make a difference without driving the less visible and more valuable ISIS supporters out of the social network in large numbers.

If every single ISIS supporter disappeared from Twitter tomorrow, it would represent a staggering loss of intelligence—assuming that intelligence is in fact being mined effectively by someone somewhere. However, many thousands of accounts can likely be removed from the ecosystem without having a dramatic negative impact on the potential intelligence yield.

4.2 Effectiveness of Suspensions in Limiting ISIS's Influence

As discussed in section 2.5, the data we collected did not provide an ideal point of comparison to evaluate suspensions. In part, this was due to an aggressive campaign of suspensions that began at the same time that we collected our initial data; this was further complicated by the weeks or months of collection time required to create an ideal comparison group.

Nevertheless, the limited data we were able to collect (section 2.5.4), as well as observational data, pointed to dramatic limits on ISIS activity online as a result of the ramped-up suspension regimen.

Data collected at various times since September 2014 consistently demonstrates that more than 8 percent of online activity by ISIS supporters is now being dedicated to rebuilding the network (sections 1.10 and 2.5.4). This figure measures only the promotion of new accounts and does not include abstract discussion of the suspensions, or discussions of strategies to compensate for the negative effects of suspensions, such as the loss of intelligence data. Based on observations, we are confident such content—combined with the promotional tweets—now constitutes more than 10 percent of ISIS supporters' online activity.

Additionally, we observed drops in activity such as replies and retweets within the network, although such data comes with significant caveats. ISIS supporters themselves also characterized the effects of the suspensions as “devastating” in strategy documents, and repeatedly emphasized the importance of creating new accounts.

Despite this, the number of new accounts created dropped significantly after the first round of suspensions in September (section 2.5.4), and while we do not have complete data to make a positive assessment, it appears the pace of account creation has lagged behind the pace of suspensions.

Most metrics within the network remained relatively flat for users who were not suspended. Despite the considerable discussion around the issue, the Demo-

graphics Dataset saw a suspension rate of only 3.4 percent between fall 2014 and January 2015. Based on the discrepancy between the suspension rate and the reports and discussions of suspensions in the content tweeted by supporters, we infer that many suspensions have targeted accounts created by users who were returning after a previous suspension. We also saw some moves to suspend ISIS-supporting bots and other accounts involved in manipulative activity.

Perhaps most important is what we didn't see. We did not see images of beheaded hostages flooding unrelated hashtags or turning up in unrelated search results. We also did not see ISIS hashtags trend or aggregate widely.

The primary ISIS hashtag—the group's name in Arabic—went from routinely registering in 40,000 tweets per day or more in around the time suspensions began in September 2014, to less than 5,000 on a typical day in February. Many of those tweets consisting of hostile messages sent by parties in the Persian Gulf.

Some stakeholders have objected to account suspensions on the basis the continuing availability of ISIS propaganda on social media and on the purported resilience of the ISIS social network.²³ These represent a straw man argument, based on the idea that the current level of suspensions is enough to destroy ISIS's presence on social media and render its content completely unavailable. No one has suggested that this is the case. This argument is the equivalent of saying we should not arrest criminals, because crime keeps coming back. While there are many useful debates in society about the level and nature of law enforcement, but there exists no legitimate argument in favor of abandoning law enforcement altogether because it's just "whack-a-mole."

There are many potential alternative outcomes short of the total eradication of ISIS-affiliated accounts, and significantly degraded performance is certainly one of them. Specifically, neutering ISIS's ability to

use Twitter to broadcast its message outside of its core audience has numerous potential benefits in reducing the organization's ability to manipulate public opinion and attract new recruits.

The data we collected also suggests that the current rate of suspensions has also limited the ISIS network's ability to grow and spread, a consideration almost universally ignored by critics of suspension tactics. The consequences of neglecting to weed a garden are obvious, even though weeds will always return.

Our data indicates that progress has been made toward limited and realistic goals at the current level of suspensions, which remains very low in comparison to the overall ISIS presence on Twitter (3.4 percent of the Demographics Dataset, not counting users who have successfully created new accounts).

Twitter has massive computing resources at its disposal, as well as access to user data at a much faster rate than permitted to outsiders. Using the techniques outlined in this paper, it is highly likely that Twitter could—if it so chose—substantially deny the use of its service to ISIS supporters, reducing their ranks to as few as a couple hundred hyper-committed supporters with negligible influence. For many reasons—including issues associated with establishing a broad precedent on political speech and the practical intelligence concerns outlined in section 4.1—we do not recommend this approach. However, it remains theoretically possible.

It is also possible to fine tune the current suspensions efforts to further limit Twitter's utility to ISIS, without completely eliminating the group's presence. For instance, given the large number of small accounts in the system, we believe it would be possible to design metrics (following established ideas such as betweenness centrality,²⁴ which describes the how a user in the network bridges between other users) that could be used to dismantle the network by separating these small accounts into ever smaller clusters of users, and disrupting the flow of information among them.

23. Ali Fisher and Nico Prucha, "ISIS Is Winning the Online Jihad Against the West," *Daily Beast*, 1 October 2014, <http://www.thedailybeast.com/articles/2014/10/01/isis-is-winning-the-online-jihad-against-the-west.html>; Jamie Bartlett and Ali Fisher, "How to beat the media mujahideen," *Demos Quarterly*, 5 February 2015, <http://quarterly.demos.co.uk/article/issue-5/how-to-beat-the-media-mujahideen/>.

24. "Betweenness Centrality," NetworkX documentation, http://networkx.lanl.gov/reference/generated/networkx.algorithms.betweenness_betweenness_centrality.html.

While Twitter does not publicly discuss its rationale for suspensions, it can be deduced that it primarily suspends accounts based on the content of tweets and user reports. This approach tends to favor accounts that are more visible and more active, consistent with our findings. Twitter could choose to approach the issue proactively with an eye toward dismantling the network rather than putting out fires as they spring up.

There are obvious ethical dimensions when a private company decides to tackle a thematic network; this research cannot inform these discussions. To some extent, regulating ISIS per se presents very few ethical dilemmas, given its extreme violence and deliberate manipulation of social media techniques. However, the decision to limit the reach of one organization in this manner creates a precedent, and in future cases, the lines will almost certainly be less clear and bright.

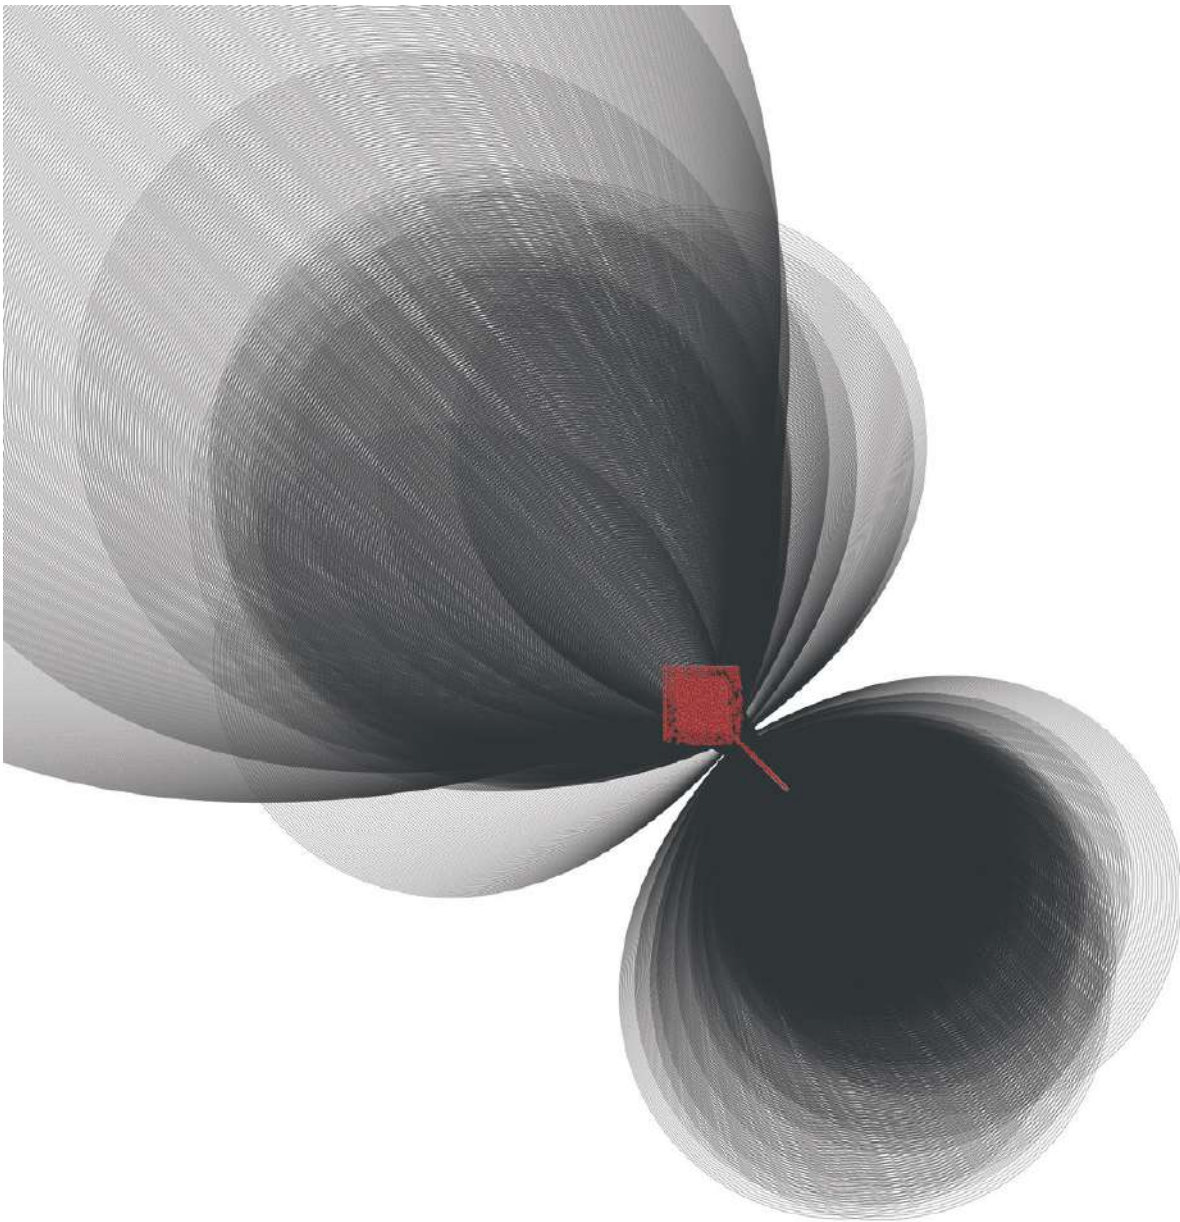


FIGURE 27: A network graph of interactions among members of the February 2015 collection set, reflecting the impact of months of suspensions. The lower portion of the graph shows a high density of interactions within the most-connected members of the set; while peripheral members in the upper section are much less connected. As suspensions contract the network, members increasingly talk to each other rather than to outsiders.

4.3 Suspensions and Trade-offs

While the current debate over the value of account suspensions leaves much to be desired, our research identified one unintended consequence that merits serious consideration as an argument against maintaining a program of suspensions. One effect of suspending more visible ISIS accounts is that the network becomes more internally focused over time. While our data was insufficient to fully address this question, the process of trying to create the February 2015 comparison set described in section 2.5.4 vividly illustrated this change in dynamic. We made several attempts to develop a network of similar size to the original collection with new seeds, only to be met with frustration. There was substantial overlap in the accounts followed by the new seed accounts, pointing toward a much more inwardly focused network, even as the average number of followers increased.

Much of what follows is observational, although future research may be able to quantify the effects. We suspect that as the network comes under mounting pressure, ISIS supporters will increasingly follow and interact with other supporters, and will be less and less inclined to follow and interact with people outside the supporter network. The repeated suspension of the most visible accounts has the effect of peeling away many users who are less engaged with ISIS and its ideology. In short, ISIS social networks on Twitter are becoming even more insular than they were to begin with.

There are potential hazards here.

First, ISIS supporters who have more out-of-network relationships may be exposed to moderating or deradicalizing influences. While ISIS's status as the most extreme Islamic radical group raises very legitimate doubts about whether the majority of

adherents are even vulnerable to moderating influences, we have seen examples of people turning away from its toxic ideology.²⁵ When we segregate members of ISIS social networks, we are, to some extent, also closing off potential exit ramps.

Secondly, while the suspensions raise the barrier to joining the social network—in the sense that they reduce the number of invitations ISIS can successfully broadcast—they do not by any means make joining impossible. The interior of this network is changing as a result of the suspensions, making it a much louder echo chamber.

The increased stridency and monotonic content may discourage some new members of the network from remaining. For others, there is a risk that the more focused and coherent group dynamic could speed and intensify the radicalization process.

Prior to the advent of social media, al-Qa'ida training camps practiced cult-like techniques of indoctrination which included cutting new recruits off from the outside world. While the barriers to outside engagement in a virtual environment are obviously far more porous, the segregation of ISIS's social network may create a smaller but similar effect. In some ways, this effect may be more pernicious as it creates the illusion of access to unfiltered information, when in fact ISIS news sources are extraordinarily filtered and biased. This allows ISIS to powerfully manipulate the selection of information its adherents can access.

This concern, however, must be weighed against other ISIS strategies, most notably its effort to encourage so-called “lone wolf” attacks by people who are only marginally engaged with its ideology and political cause.²⁶ Some lone actors—such as those who carried out attacks in Canada in October²⁷—have been fully engaged with ISIS social networks. Others—such as the man who attacked

25. “ISIS defector on why he left extremist group, how he escaped,” *CBS News*, 11 February 2015, <http://www.cbsnews.com/videos/isis-defector-on-why-he-left-extremist-group-how-he-escaped/>.

26. J.M. Berger, “The Islamic State’s Irregulars,” *Foreign Policy*, 23 December 2014, <http://foreignpolicy.com/2014/12/23/the-islamic-states-deranged-irregulars-lone-wolf-terrorists-isis/>.

27. “Suspect in Quebec attack on soldiers was arrested in July: RCMP,” *CTV News*, 21 October 2014, <http://www.ctvnews.ca/canada/suspect-in-quebec-attack-on-soldiers-was-arrested-in-july-rcmp-1.2063243>; Saeed Ahmed and Greg Botelho, “Who is Michael Zehaf-Bibeau, the man behind the deadly Ottawa attack?,” *CNN Online*, 23 October 2014, <http://www.cnn.com/2014/10/22/world/canada-shooter/>.

New York police with a hatchet²⁸—were less visibly engaged and may be more psychologically similar to spree killers than to traditional conceptions of networked terrorists.

Regardless of where on the spectrum such individuals lie, research indicates that mental illness plays a significant role in lone-actor terrorism, and ISIS's ultraviolent propaganda provides an unusually high level of stimulation to those who might already be prone to violence.

While it is relatively easy for someone with knowledge of Iraqi and Syrian social networks and information outlets to find such graphically violent ISIS content, the suspensions do provide some degree of buffer against mentally ill individuals who might seek out images of sadistic violence to strengthen their own violent impulses, particularly in English. Again, this is not to say that such material is unavailable, but pressure on ISIS's social networks does reduce the distribution of such material to people who do not have a specialized interest.

Fundamentally, the questions raised here pertain to the largely unacknowledged fact that tampering with social networks is a form of social engineering. There are considerations that go beyond the instrumental question of how ISIS spreads propaganda and recruits, cutting to the heart of how social radicalization works and raising questions about unintended consequences. Additional study of these issues should be a top priority.

4.4 Preliminary Policy Recommendations

No single authority possesses the scope and power to fully address the challenges presented by the presence of ISIS and other similar groups on social media. There are, however, stakeholders, and each must answer questions about their own internal priorities in this conflict. To achieve progress, stakeholders must come together in an organized

way to discuss potential solutions. Although this study has examined Twitter and ISIS in particular, the issues and dynamics laid out here also apply to other social media platforms and other groups exerting negative social impact, such as white nationalists, whose networks on Twitter have developed substantially since the previously referenced 2013 study "Who Matters Online."

Social media platforms must first ask themselves where their responsibilities lie. While we do not believe that any mainstream social media platform wishes to see its services used to further acts of horrific violence, we also suspect some would rather not be bothered with the challenge of crafting a broad and coherent response to the issue.

While we can sympathize with the challenges and dilemmas such a response would entail, it is clear that social media companies do feel an obligation to respond to some social standards and illegal uses of their services. We are not aware of any major company that takes a hands-off approach to the use of its platform to promote child pornography or human trafficking—or less dramatically, phishing, spam, fraud, and copyright violations. Extremism, while raising thornier issues, merits attention, especially when faced with a rising challenge of violent groups who manipulate platforms to reap the rewards of spreading images of their cruelty. Some platforms—notably Facebook and YouTube—have already instituted policy changes specific to extremism.

Social media platforms should consider whether they want to continue with some variation of their current approach, which tends to stomp out fires as they erupt, or whether they want to dismantle or degrade the social networks responsible for setting the fires. While we do not necessarily recommend that social media platforms take the network-wide approach, it should nonetheless be examined and considered in greater depth using social network analysis techniques such as those featured in this study.

28. James Gordon Meek and Josh Margolin, "NYC Ax Attacker Was Consumed by Desire to Strike U.S. Authority Figures, Police Say," *ABC News*, 3 November 2014, <http://abcnews.go.com/US/nyc-ax-attacker-consumed-desire-strike-us-authority/story?id=26664787>.

By understanding the power and the limits of a holistic approach to network degradation, we can better understand what approaches might be most effective and least intrusive. Such study should also examine in depth how isolating extremist communities may have counterproductive effects.

Social media platforms have thus far been able to take shelter in a presumption that their platforms are protected as Internet service providers under laws written to exempt the phone company from liability for illegal acts carried out using a phone (such as using a phone to arrange a purchase of drugs). Given that social media platforms have the capacity to broadcast messages akin to a television or radio broadcast, the law may eventually change to keep up with the evolution of technology.

Nevertheless, even telephones are regulated for antisocial content. For instance, telephone harassment is illegal in all 50 states,²⁹ but social media harassment is not, despite the fact that using social media to transmit pictures of gory executions can be more intrusive and pernicious than some behaviors banned under telephone statutes.

It is unwise for social media companies to presume they will remain immune to regulation. Companies should get out ahead of the curve by crafting policies and publicly articulating their priorities. If they do not bring their vision to the government, the government is likely to bring a much more restrictive vision to them.

Government, for its part, must do something it has not traditionally excelled at: fully address a complex situation and attempt to find a nuanced approach. Most obviously, social media is embedded with a component of free speech, and in recent years has drawn attention to international crises and fostered popular dissent in authoritarian societies.

Any attempt to establish sweeping authority over political speech on social media comes with po-

tentially high risks, and laws tuned specifically to ISIS are problematic on multiple levels. Given that the major social media platforms are based in the United States, the U.S. government has the most obvious authority over how these companies are managed. Nonetheless, the government also has an obligation—embedded in our constitutional principles—to protect users and the companies themselves from authoritarian abuse.

Discussions of the regulation of speech on social media platforms tend to emphasize libertarian values and the protection of free speech from government intrusion. However, the legal vacuum that currently surrounds these issues is a *de facto* concession of near-absolute authority to corporations, rather than the empowerment of users.

This point needs to be crystal clear: social media companies can and do control speech on their platforms. No user of a mainstream social media service enjoys an environment of complete freedom. Instead, companies apply a wide range of conditions limiting speech, using possibly opaque guidelines that may result in decisions executed on an *ad hoc* basis. Furthermore, companies typically do not disclose information about who they suspend and why, nor are they required to.

There are many attendant questions that should be of interest to civil libertarians. These include, for example, whether suspensions disproportionately impact people of certain genders, races, nationalities, sexual orientations, or religions. Twitter in particular discloses literally no information about the accounts it suspends, yet this activity takes place every day. Again, this is an area in which companies would be well-advised to consider proactive measures, and it is an area where government oversight may eventually come into play.

It is apparent that progress will only be made when these two largest stakeholders—private sector providers and the government—come to-

29. Josh Tatum, "Telephone-harassment statutes," First Amendment Center, 18 August 2008, <http://www.firstamendmentcenter.org/telephone-harassment-statutes>.

gether and discuss the many legitimate concerns on both sides.

This process should include input from other stakeholders as well, including experts in extremism and violent extremism, victims of extremist violence, activists working to counter extremism, scholars of political dissent, and advocates of civil liberties. Any one-sided or even binary solution—conceived of and imposed primarily by some combination of government and corporations—will likely to be inadequate to deal with the complex issues raised by the problem of extremist use of social media.

4.5 Research recommendations

We close by emphasizing our belief that while this research significantly advances the state of knowledge about the functioning of extremist social networks, it does not present a complete picture of the effects of suspending social media accounts used by extremists. We faced two significant challenges, which could be remedied in future studies using similar methodologies.

First, due to the timing of the study, we were unable to create a dataset that could serve as a control group for how ISIS functions in a low-suspension environment. Aggressive suspensions started at the same time that we began collecting data.

The most likely implication of this limitation is that we have understated the impact of aggressive suspensions in degrading the performance of ISIS social networks. Various data points collected opportunistically by J.M. Berger since 2013, such as hashtag trends, tend to support this view.

It may be possible to draw on commercially available Twitter data to create a historical comparison dataset that would provide a certain degree of improved resolution on this question. Ultimately, based on the limitations of the cur-

rent commercial offerings, the only approach to creating a comprehensive comparison would require the disclosure of data held internally by Twitter, a prospect that seems unlikely.

Second, this research clearly points to possible methodologies for actively monitoring the evolution of ISIS's social network over time, allowing for much more precise evaluations of the number of suspensions occurring over time, and the short and long term effects of those suspensions.

A project of this sort would create a massive reservoir of open-source data on ISIS and its use of social media, and it is technically feasible barring a sea change in Twitter's rules for data access.

Additional study would provide a much clearer view on how the network responds to suspensions, how it evolves, and how its internal social dynamics change over time. We hope that the detailed disclosure of the methodology used to develop this report and its utility for future research will advance the study of social network dynamics in general, and extremist use of social media in particular.

In principle, the best approach this problem would be to run an analysis similar to that used in this study on an iterative and continuing basis—after addressing some technical challenges, including the size of the dataset and the amount of time required for collection. Analysts could then define new seed accounts from the collected data, according to replicable criteria, run the complete analysis, and then create new seeds using the same criteria and run the analysis again.

The first iteration of this process would likely create a meaningfully different data set, due to the original seeds being chosen using manual selection by a human analyst; subsequent iterations would likely become more uniform and more directly comparable, with changes to the dataset reflecting the organic evolution of the network.

Appendices

A. Notes on In-Network Calculations

User information is available more quickly via the Twitter application programming interface (API) than information about friends and followers, which is considerably slower to collect. The rate at which friend and follower information can be collected is significantly limited by Twitter and cannot be negotiated.

This proved relevant to the final multiplier in the metric, in which the number of out-network friends and followers were computed using the database record for “total friends” and “total followers.” We collected more friends than the “total friends” entry for 7.3 percent of users. For followers, the figure was 1.6 percent.

The discrepancy was lower for followers because the in-network followers was collected from user information, which was weighted toward identifying friend relationships. In-network followers were computed by looking at how many people within the network had listed the user as a friend.

This introduced some bias into the metrics due to the fact that collection was spaced out over time. In other words, if a user was collected early in the process, the discrepancy between the database entry and the in-network calculations would tend to be lower; if the user was collected later in the process, it would tend to be higher.

Having not fully anticipated this issue or the importance of in/out metrics in advance of collection, we were unable to account for the discrepancy manually. If we had anticipated it, we could have made structural changes to the collection process to partially compensate. By the time the significance of the discrepancy became clear, the challenges of trying to regressively account for the change outweighed the benefits.

This was in part because the discrepancy ultimately proved advantageous. When an account added a substantial number of in-network relationships in the intervening time between collection of user information and collection of relationships, the velocity itself became a partial indicator of the relevance of the account: in other words, highly relevant accounts added relationships faster than less relevant ones. Since the metric used several other characteristics to evaluate responsiveness, other data was available to offset the possible negative consequences while still gaining some of the benefits. Some inherent bias related to the timing of collection is likely unavoidable under Twitter’s current terms of service for API access.

In future projects, by estimating the velocity at which in-network friends were added, it might be possible to improve the accuracy of the process of sorting accounts. However, this approach—and the associated caveats—are almost certainly more applicable to ISIS supporters than to general interest queries, due to the highly organized manner in which ISIS builds its network, and the rather specific behaviors related to the aggressive suspension of supporter accounts.

B. Control Group

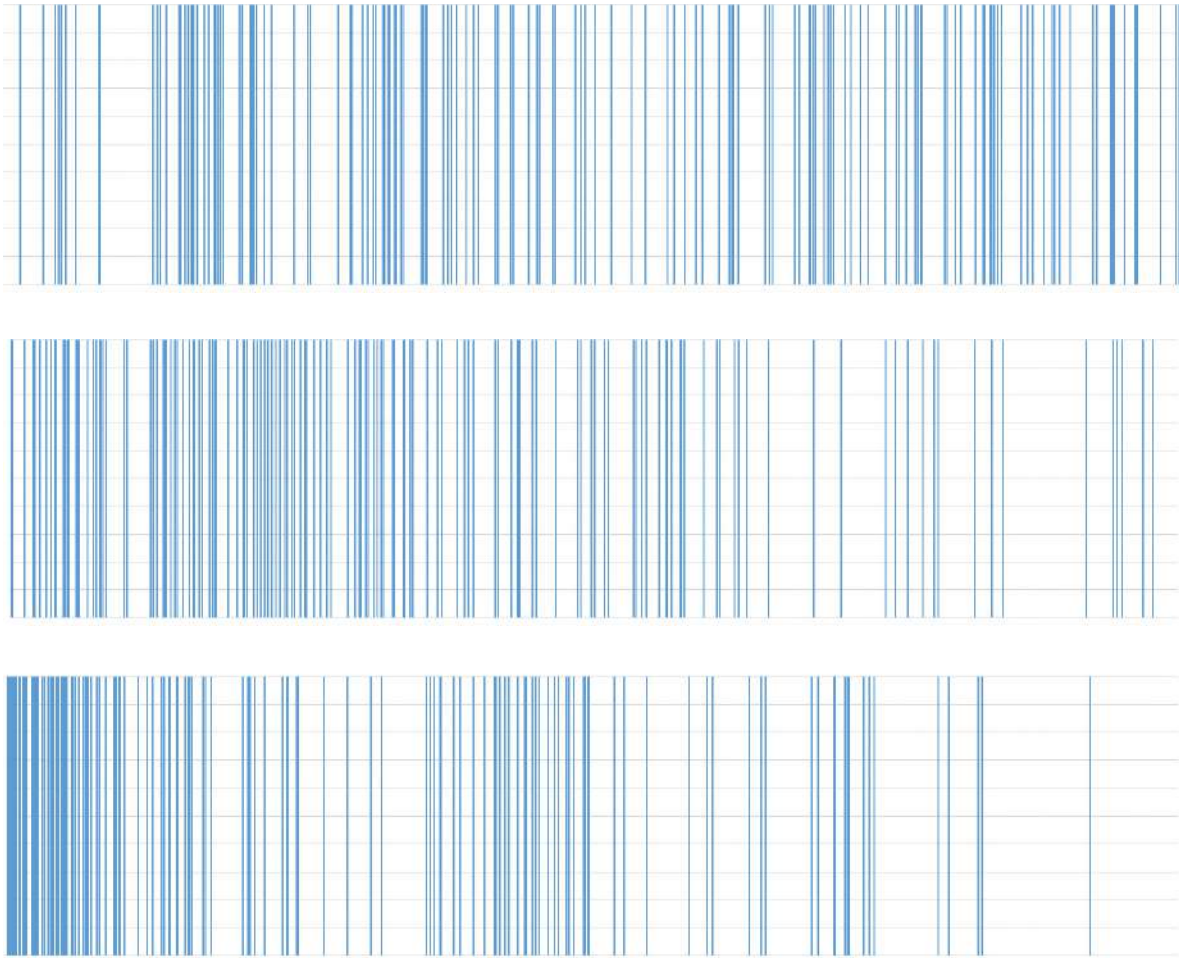


FIGURE 28: The data journalist set sorted 1) alphabetically, 2) by engagement, 3) using the IQI metric

We tested the sorting metric on a non-extremist dataset, using four prominent data journalists as seed accounts, to determine whether it would effectively sort data journalists (a much more specific subset than ISIS supporters) and people interested in data journalism from the larger social network of accounts followed by the seed users. The object was to provide a comparison group to the ISIS analysis, and also to explore the utility of the metric outside of extremist networks.

For a quick evaluation, we coded as responsive hits accounts with a profile description that included the word data combined with one of the following terms: journalist, journalism and reporter. The IQI metric (without the adjustment for non-reply mentions) again proved considerably more effective

than other metrics, although slightly less effective than it was in sorting ISIS supporters.

Extremist groups are by nature insular. Members prefer to get news and information from other members of the group, and their desire for interactions tends to be contained within the group (ISIS's external broadcast of propaganda as a notable exception).

We attribute the difference in effectiveness to this element, although it may also be related to the fact that the data journalist set was considerably smaller than the ISIS set. We suspect additional data would contribute to a more accurate sorting process, likely by weighting lower-ranked responsive accounts more accurately through the collection of more members of their networks.

About the Project on U.S. Relations with the Islamic World

The Brookings Project on U.S. Relations with the Islamic World is a research initiative housed in the Center for Middle East Policy at the Brookings Institution. The Project's mission is to engage and inform policymakers, practitioners, and the broader public on the changing dynamics in Muslim-majority countries and to advance relations between Americans and Muslim societies around the world.

To fulfill this mission, the Project sponsors a range of activities, research projects, and publications designed to educate, encourage frank dialogue, and build positive partnerships between the United States and Muslim communities all over the world. The broader goals of the Project include:

- Exploring the multi-faceted nature of the United States' relationship with Muslim-majority states, including issues related to mutual misperceptions;
- Analyzing the social, economic, and political dynamics underway in Muslim societies;
- Identifying areas for shared endeavors between the United States and Muslim communities around the world on issues of common concern.

To achieve these goals, the Project has several interlocking components:

- The U.S.-Islamic World Forum, which brings together leaders in politics, business, media, academia, and civil society from the United States and from Muslim societies in Africa, Asia, Europe, and the Middle East. The Forum also serves as a focal point for the Project's ongoing research and initiatives, providing the foundation for a range of complementary activities designed to enhance dialogue and impact;
- An Analysis Paper Series that provides high-quality research and publications on key questions facing Muslim states and communities;
- Workshops, symposia, and public and private discussions with government officials and other key stakeholders focused on critical issues affecting the relationship;
- Special initiatives in targeted areas of demand. In the past these have included Arts and Culture, Science and Technology, and Religion and Diplomacy.

The Project's Steering Committee consists of Martin Indyk, Vice President and Director of Foreign Policy Studies; Tamara Wittes, Senior Fellow and Director of the Center for Middle East Policy; William McCants, Fellow and Director of the Project on U.S. Relations with the Islamic World; Kenneth Pollack, Senior Fellow in the Center; Bruce Riedel, Senior Fellow in the Center; Shibley Telhami, Nonresident Senior Fellow of the Project and Anwar Sadat Chair for Peace and Development at the University of Maryland; and Salman Shaikh, Fellow and Director of the Brookings Doha Center.

The Center for Middle East Policy

Charting the path to a Middle East at peace with itself and the world

Today's dramatic, dynamic and often violent Middle East presents unprecedented challenges for global security and United States foreign policy. Understanding and addressing these challenges is the work of the Center for Middle East Policy at Brookings. Founded in 2002, the Center for Middle East Policy brings together the most experienced policy minds working on the region, and provides policymakers and the public with objective, in-depth and timely research and analysis. Our mission is to chart the path—political, economic and social—to a Middle East at peace with itself and the world.

Research now underway in the Center includes:

- Preserving the Prospects for Two States
- U.S. Strategy for a Changing Middle East
- Politics and Security in the Persian Gulf
- Iran's Five Alternative Futures
- The Future of Counterterrorism
- Energy Security and Conflict in the Middle East

The Center was established on May 13, 2002 with an inaugural address by His Majesty King Abdullah II of Jordan. The Center is part of the Foreign Policy Studies Program at Brookings and upholds the Brookings values of Quality, Independence, and Impact. The Center is also home to the *Project on U.S. Relations with the Islamic World*, which convenes a major international conference and a range of activities each year to foster frank dialogue and build positive partnerships between the United States and Muslim communities around the world. The Center also houses the *Brookings Doha Center* in Doha, Qatar—home to three permanent scholars, visiting fellows, and a full range of policy-relevant conferences and meetings.

BROOKINGS

The Brookings Institution
1775 Massachusetts Ave., NW
Washington, D.C. 20036-2103

www.brookings.edu

