




# Defeating ISIS on Twitter

Batsheva Moriarty

## Highlights

- I evaluated 1.5 million tweets from 1,500 ISIS-affiliated Twitter accounts to determine if they were humans or bots
- I compared ISIS tweets to a control group of 700,000 non-ISIS Arabic tweets
- ISIS tweets exhibited unique, un-unified tweet, retweet, and favoriting patterns suggesting that the accounts are controlled by humans



*Proportion of retweet rates of suspected ISIS-supporters versus control group.*

## Abstract

ISIS, also known as ISIL, the Islamic State, and Daesh, is a Sunni terrorist organization that has seized large portions of Iraq and Syria [1]. An estimated 20,000 fighters from 90 countries, including as many as 150 Americans, have traveled to Iraq and Syria to join the ranks of ISIS [2]. In addition to military might, ISIS uses social media to propagate its ideas and recruit members. According to a study by the Brookings Institution, ISIS supporters used at least 46,000 Twitter accounts between September and December 2014 [3]. In this experiment, I

explore whether these accounts are likely to be human- or machine-powered. Specifically, I study the way ISIS-supporting accounts tweet, retweet, and favor messages.

**Results summary:** I examined 1.5 million tweets posted between August 2014 and April 2015 by 1,500 Twitter accounts reportedly associated with ISIS, and compared them to 700,000 tweets from a control group of comparable Arabic-language accounts not believed to be associated with ISIS. I compared the frequencies of tweets, retweets, and favorited tweets. ISIS accounts did not seem to tweet in a unified manner, had fewer retweets than the control group, and spikes in ISIS tweets associated more strongly with regionally important ISIS events than with worldwide ISIS events. These findings suggest that primarily humans, and not machines (or bots), tweet and retweet from Twitter accounts that are reportedly associated with ISIS.

## Introduction

According to a February Gallup poll, 84 percent of Americans consider ISIS a critical threat to the United States [4]. ISIS' stated goal is to restore the caliphate, an Islamic state guided by a supreme leader [5]. The group is notorious for its brutality and for enslaving, raping, and executing Christian, Yazidi, and Western captives.

According to FBI Director James Comey, "Homegrown violent extremists are troubled souls who are seeking meaning in some misguided way, and so they come across the propaganda, and they become radicalized on their own" [6]. In some cases, the radicalized behavior may result in violence within their home countries. Once Westerners began leaving their home countries to join ISIS after interacting with it on social media, international focus moved to ISIS's social-media campaign.

Research on ISIS' social-media accounts has been limited, with one study by the Brookings Institution providing the vast majority of empirical information on the topic [3]. The Brookings study found that ISIS uses social media strategically to disseminate images of graphic violence, attract new recruits, and potentially inspire lone-actor attacks [3]. To dissuade people from supporting ISIS, the U.S. State Department launched its own social-media efforts, including a controversial Twitter counter-messaging campaign, to contest the extremists' claims on social media [7]. In cooperation with the State Department, Twitter shut down more than 10,000 ISIS-linked accounts in April 2015 [8].

Understanding ISIS' social-media strategy, including who runs the accounts and whether they operate in an organized manner, would enable the American government, with help from technology companies, to work out a more effective strategy of combatting ISIS propaganda. Disrupting its effectiveness on social media may result in fewer recruits joining ISIS, and potentially inhibit lone-actor attacks. A critical question is whether thousands of humans are executing ISIS' social-media campaign, or whether the group is using automated systems to echo messages sent by a few people.

## Background

There are many ways to use technology to tweet or retweet messages. These include browsers and automated scripts running on computers, apps operating on mobile devices, and professional services.

The Brookings Study identified apps commonly used by ISIS supporters [3]. Many of these apps were not designed specifically for ISIS' use. For example, BufferApp schedules tweets to be sent at a particular time. Hootsuite and Tweetbot simply allow users to tweet and follow others.

According to the Brookings study, the app Dawn of Good Tidings was endorsed by top ISIS online personalities, and at its peak sent tens of thousands of tweets per day [3]. Twitter terminated the app in June 2014. Reportedly, ISIS supporters responded to the shutdown 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 wished to promote. If one bot family gets suspended, many others continue to tweet [3].

Are ISIS tweets primarily human-generated, or automated?

## Methods

My approach is to collect tweets believed to be from ISIS accounts and collect a similar number of tweets from comparable accounts that are not necessarily ISIS-related. I then examine and compare the frequencies of tweets, retweets, and favorites over time to spot patterns and to draw inferences.

### Data Collection

I acquired a list of operating Twitter accounts identified by the hacktivist organization Anonymous as belonging to ISIS [9]. Anonymous seemingly posted this list to encourage Twitter to take them down [9]. I could not verify that these accounts were actually related to ISIS, but based on a manual examination of the English-language content, they appear to belong to ISIS. For this study, I refer to these as the ISIS accounts and the tweets recovered from them as ISIS data.

The Twitter API rate limit, which allows only 720 API calls per hour and returns a maximum of 3,200 tweets per account, was a major limitation to data collection, so I decided to focus on the accounts that would have the most useful data. I eliminated accounts with fewer than 100 followers, because those likely would not play a major role in ISIS' social-media strategy. I also eliminated accounts that were less than six months old because those would not allow me to look at trends over time. Lastly, I eliminated accounts that tweeted, based on annualized monthly rates, more than 7,000 times per year, because the maximum of 3,200 tweets that I could retrieve likely would not span enough months. These eliminations left me

with approximately 1,500 accounts and 1.5 million tweets spanning August 2014 to April 2015. I collected these tweets at the beginning of May 2015. For each tweet, I collected the time of the post, the number of times it was retweeted, the number of times it was favorited, the language of the tweet, and the text of the tweet if the language was English. I also saved the account name and its number of followers.

### Establishing a Control Group

In order to draw accurate conclusions about the data, I needed a comparable control group of regular Twitter users. The vast majority of the tweets in my initial dataset were in Arabic, so to get a group with a comparable demographic I collected the usernames of random tweeters who were posting in Arabic. I used the Twitter Stream API to collect usernames, and then filtered for accounts that tweeted less than 7,000 times per year and satisfied other criteria I used for ISIS accounts. I then used the same script I used in my initial data collection to collect the same information about each Twitter account, such as the time of each tweet, the number of retweets, favorites, and so on. Paring down the dataset by follower count and average tweeting rate could only be done once all the data was collected, so I ended up with approximately 700,000 tweets that comprised a random comparable dataset of Twitter histories that could be collected given the limitations of the Twitter API.

### Methodology

I combined all the Twitter data from the ISIS group and the control group into one master DataFrame using the Python Pandas library. I separated tweets by month to examine frequencies over shorter periods. The data had more tweets in the later months, so examining months individually helped account for monthly variation. Because I do not know Arabic, I focused entirely on frequency analysis and did not look at content, apart from a manual examination of the English-language tweets.

For each month of data, I used a Python script to separate tweets by day, and then used the Matplotlib to graph the total tweet count, the sum of all the retweets, and the sum of all the favorites. I defined outliers as accounts that posted over 150 tweets more than their daily median amount in a single day, and printed summary information for any outliers that appeared in the data.

I also produced a list of dates of events important to ISIS and the Middle East and compared them to the timing of ISIS tweets. On September 23, 2014, President Obama announced the establishment of a large international coalition against ISIS, and the coalition launched its first strikes on militants in Syria that day [10]. The video of James Foley's beheading was released on August 19, 2014, and the video of Steven Sotloff's beheading was released on September 2, 2014. Both of those videos made major headlines in the West. The video of the burning alive of Jordanian pilot Moaz al-Kasasbeh was released on February 3, 2015. The Muslim holiday of Eid al-Adha started October 4, 2014. Therefore, I examined tweet and

retweet behaviors on and near those five dates, which I termed “dates of interest.” I expected a change in tweet frequency on these days.

In this study, I consider a Twitter account human-powered if it demonstrates autonomy in tweeting or retweeting behavior. I consider an account machine-powered or a bot if it demonstrates rigidity in tweeting or retweeting behavior. I am not asserting that the accounts are actually humans or machines, only that the behavior of the account resembles human autonomy or machine automation. Humans can actually tweet like machines and machines may be sophisticated enough to tweet as if they give content some human deliberation.

Two scenarios inspire me to compare tweets and retweets of original content. The first scenario is episodic. If a human-powered account tweets a message, then machines may automatically retweet the message. If so, I expect to see a consistent multiplier of retweets to tweets. On the other hand, if humans retweet, I expect to see variability in the number of retweets.

The other scenario is consistent broadcast. In this case, I expect machine-powered accounts to consistently tweet and retweet.

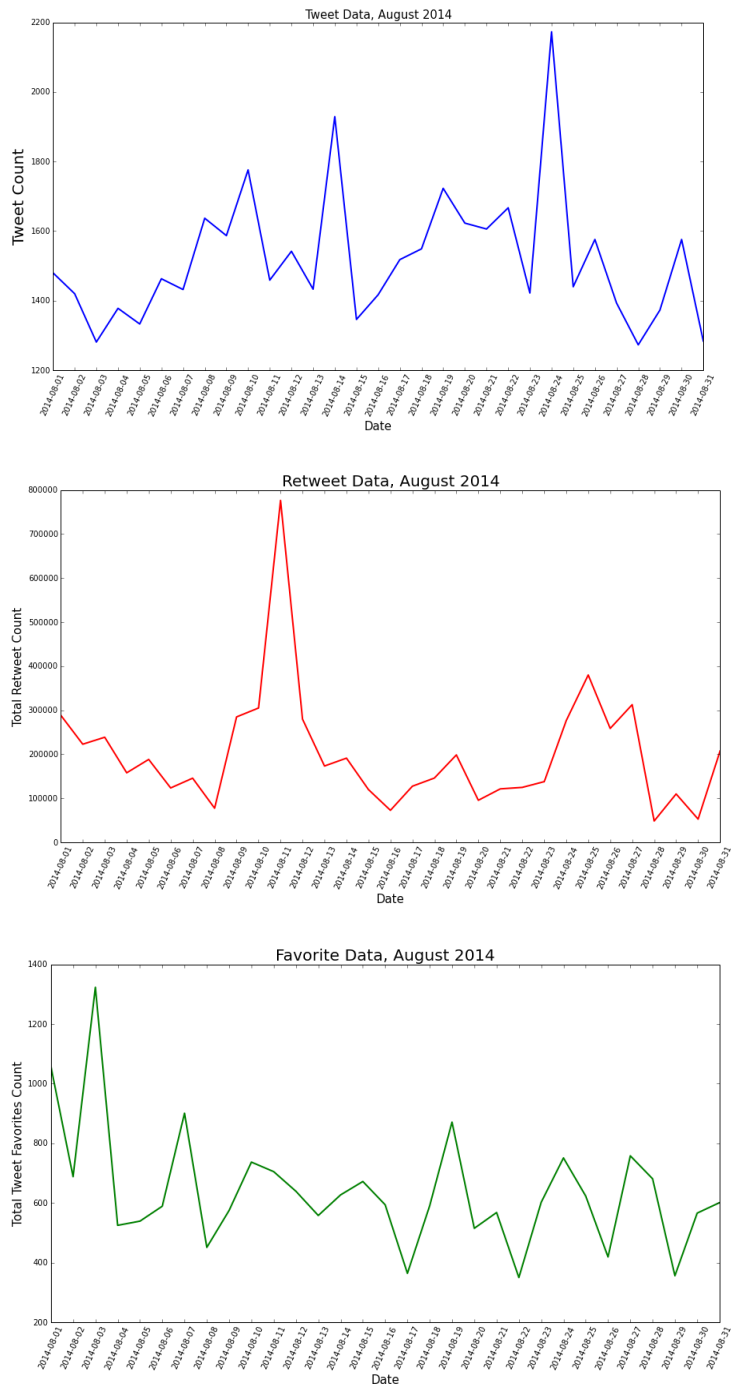
## Results

Figure 1 shows graphs of the counts of tweets, retweets, and favorites for each month from August 2014 to April 2015. In the August data, the three outliers that appear in Figure 2 contribute to the spikes in tweets on August 14 and August 24. Because these spikes are primarily caused by a few outliers, the dates themselves may not have broader relevance to explain the large number of tweets.

Figure 1 also shows that the ISIS accounts do not seem to tweet a standard amount every day, in clear trends, or in repeating patterns. Each month had widely fluctuating tweet-frequency behavior.

Figure 3 shows that in the month of January daily tweeting across 100 accounts and across 500 accounts was inconsistent, and individual accounts did not appear to follow a set pattern of tweeting. Therefore, from my definitions, I interpreted the ISIS accounts as exhibiting human behavior in tweeting.

In addition to a lack of unified or patterned tweeting among ISIS accounts, they have a lower proportion of highly retweeted posts than the control group, as shown in Figure 4. This suggests that there is a higher proportion of original content and less bot-like activity. These results also suggest that the ISIS accounts are not automatically broadcasting.




**Figure 1. Tweets believed to be related to ISIS' social-media strategy by day from August 2014. (a) Tweets. (b) Retweets. (c) Favorite tweets.**

Potential Outlier: 363FFFFF  
The date is: 2014-08-24  
They tweeted 318 times

Potential Outlier: ali7001626  
The date is: 2014-08-14  
They tweeted 421 times

Potential Outlier: al\_mki2  
The date is: 2014-08-24  
They tweeted 195 times

Figure 2. Three outliers found among August tweets from accounts believed to be related to ISIS. Each of these accounts posted over 150 tweets more than their own daily median tweet amount on the date listed.



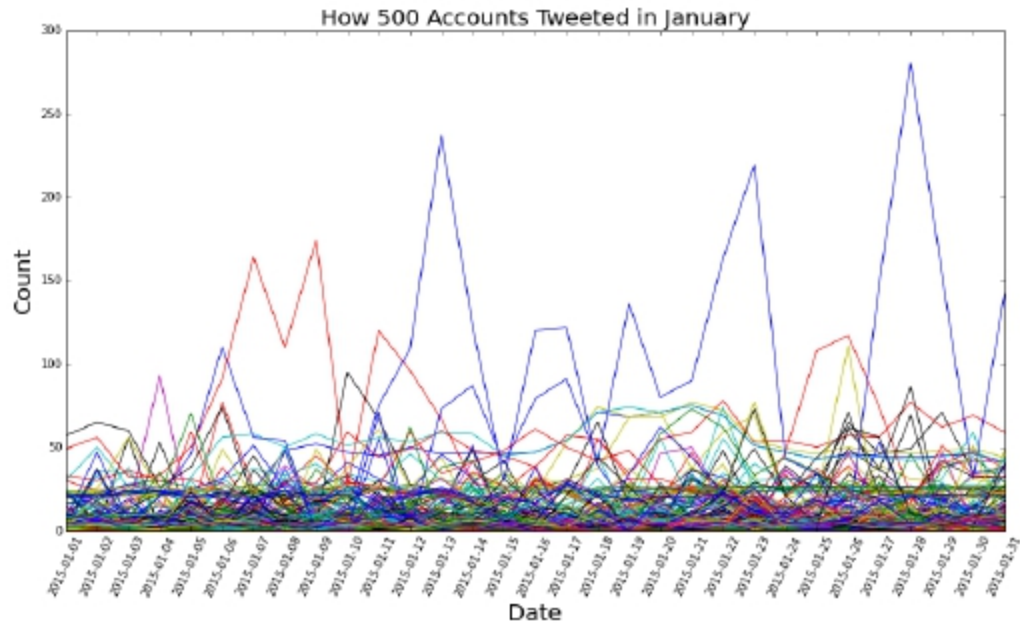


Figure 3. Number of tweets by (a) 100 accounts and (b) 500 accounts in January 2015 from accounts believed to be related to ISIS.

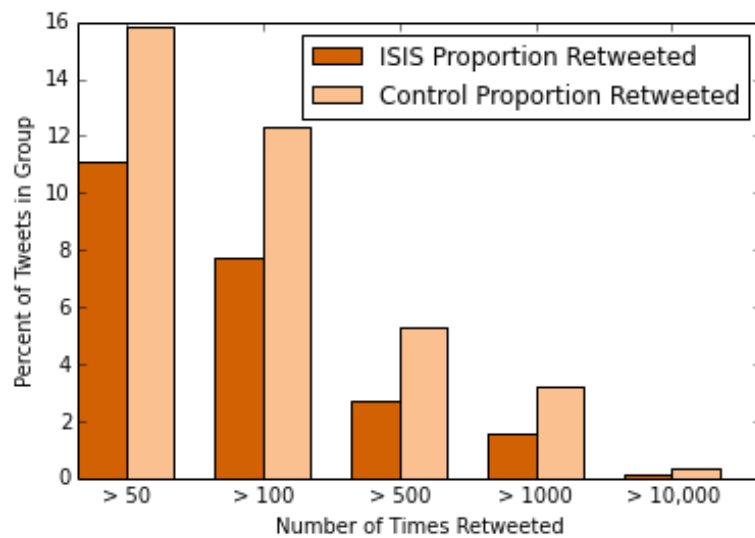
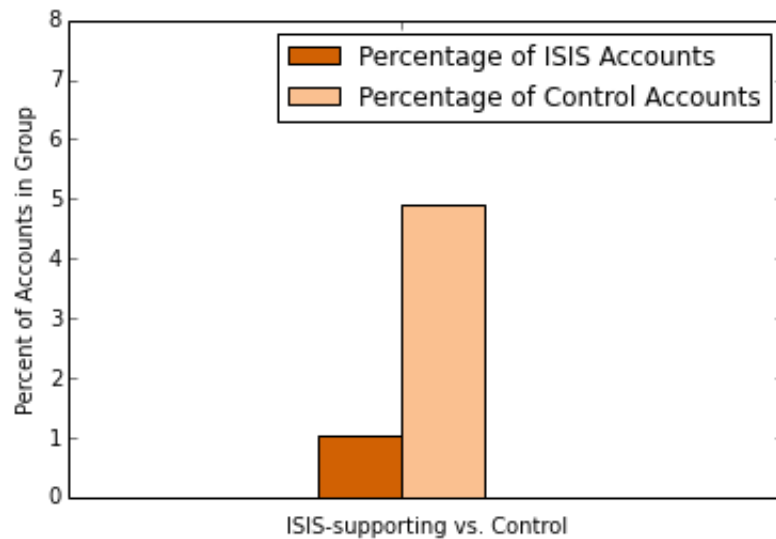


Figure 4. Proportion of retweets in group: ISIS-supporting vs. Control.






**Figure 5. Percentage of accounts in group with greater than 50 percent of tweets retweeted more than 100 times: ISIS-supporting vs. Control.**

Figure 4 compares the percentage of tweets in the ISIS accounts to those in the control group that were retweeted more than a given number of times. Twelve percent of control-group posts were retweeted more than 100 times, whereas less than 8 percent of ISIS-supporting posts were retweeted more than 100 times. Figure 5 shows that 5 percent of the control-group accounts have tweet histories where at least half of their tweets were retweeted more than 100 times. This suggests that either these accounts are being retweeted constantly or they are retweeting popular tweets constantly. Only 1 percent of ISIS accounts exhibit this behavior. If there were an ISIS strategy to propagate specific tweets, presumably there would be the same or a larger proportion of ISIS accounts devoted to retweeting. Therefore, I interpret the ISIS accounts as being human-powered.

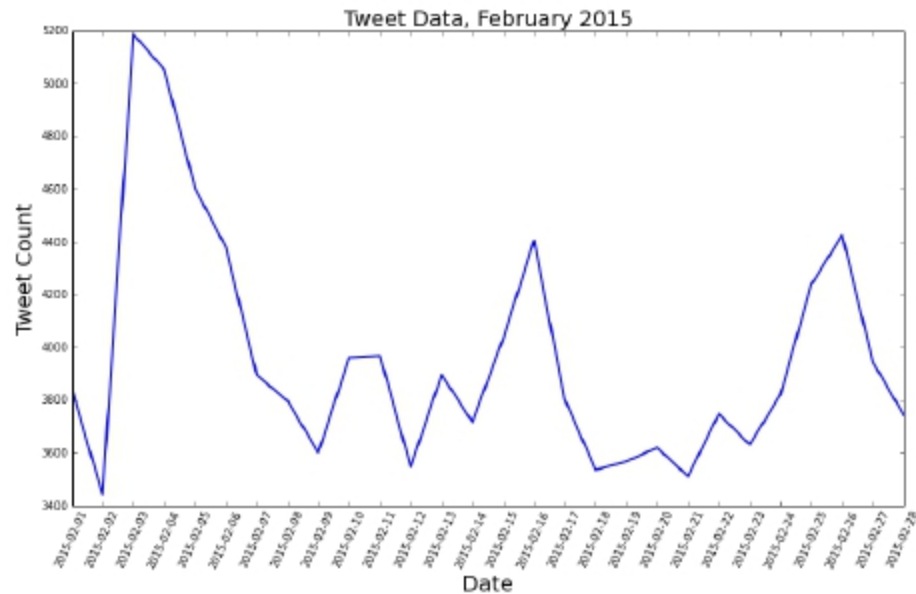
By visual inspection, I found a number of key associations in tweet frequency on the dates of interest to ISIS and the Middle East. In August, the two largest spikes in tweet count seemed to be caused by single outliers (accounts that posted over 150 tweets more than their own daily median tweet amount on the date). However, in other months there were large spikes not caused by outliers. For instance, there was a major spike in tweets, retweets, and favorites on September 23, a date of interest. Both the coalition and the strikes against ISIS may have contributed to a surge in tweets on that day. The graph of tweet data for September appears in Figure 6.



**Figure 6. Tweet counts for September 2014.**

There was no significant fluctuation in tweet or retweet frequencies surrounding the Foley and Sotloff videos that made big headlines in the West. However, the release of the video of the Jordanian pilot being burned alive did correspond to February 3 and February 4 having the largest number of tweets in the month of February, as illustrated in Figure 7. Finally, on Eid al-Adha, a Muslim holiday starting October 4, there was a noticeable dip in tweets.

The tweet peaks in apparent response to the formation of the coalition and to the murder of al-Kasasbeh and the lack of change in tweet frequency in response to events that made headlines in the West suggest that these ISIS-supporting accounts may be reacting organically to events around them. Some of these accounts may support ISIS from locations in Jordan or other Middle Eastern countries. Both the forming of the coalition, which includes various Gulf States, and the murder of the al-Kasasbeh would likely be of greater importance to ISIS supporters in the Middle East than the murders of Western journalists. It is also possible that these Twitter accounts are part of a concerted effort to reach out to Arabic-speaking audiences, and that ISIS feels that those audiences would be more interested in the murder of a Jordanian pilot than in the murders of Westerners.



**Figure 7. Tweet data February 2015. Peaks and valleys may indicate an ISIS response to events.**

Perhaps most suggestive of the role of humans in these Twitter accounts is the drop in tweets on the Muslim holiday. ISIS bots might increase tweeting in response to relevant world events, but it seems less likely that bots would decrease tweeting frequency for a religious holiday.

## Discussion

ISIS tweets appear to be driven primarily by human activity and not by bot activity nor by humans retweeting bots. First, the ISIS accounts do not seem to tweet in a unified manner. Second, the dataset of ISIS tweets has a lower proportion of highly retweeted tweets than a control group of non-ISIS Arabic speakers, which suggests a small likelihood of bot activity. Third, spikes in tweets seem to correlate more strongly with regional ISIS events than with worldwide ISIS events.

There were a number of limitations to this experiment. Ten thousand ISIS accounts were suspended shortly before I collected my data, which may be one reason for less bot-like behavior among the accounts.

Collecting information over a longer timespan would provide a much larger and more reliable dataset. Streaming data with the Twitter API in real time instead of collecting historical data would also allow for an analysis of how follower counts change over time, something that might reveal the impact of ISIS's tweeting strategies as well as the effects of ISIS-related world events. Following up on the data in this way would also enable an analysis of the benefits and consequences of shutting down ISIS accounts, as the dataset would contain

information from both before and after any shutdown. A larger and more comprehensive analysis could also explore the content of tweets and account descriptions. Studying the personal descriptions of ISIS Twitter account holders could provide a better window into the makeup of the terrorist network. Furthermore, studying the content of tweets sent, received, and retweeted during spikes and dips in the data could help ascertain causes and effects.

While I was not able to conduct a statistically significant content analysis of any of the accounts due to the language barrier, I found it interesting that tweets about soccer and the World Cup appeared prominently in a number of the English-language accounts.

Whether or not ISIS's strategy still includes bots, the seeming human engagement in ISIS' tweeting behavior suggested by this study means that ISIS tweeters may increasingly become indistinguishable from non-ISIS human accounts.

## References

1. A Visual Guide to the Rise of the Islamic State. The New York Times. Updated 11 June 2014. Web. <http://www.nytimes.com/interactive/2014/06/12/world/middleeast/the-iraq-isis-conflict-in-maps-photos-and-video.html>.
2. Kukil B. ISIS Continues Steady Recruitment As 20,000 Foreign Fighters Join Extremist Groups In Syria, Iraq: Report. International Business Times. Web. February 11 2015. <http://www.ibtimes.com/isis-continues-steady-recruitment-20000-foreign-fighters-join-extremist-groups-syria-1812440>.
3. Berger J and Morgan J. The ISIS Twitter Census. The Brookings Institute. March 2015. [http://www.brookings.edu/~media/research/files/papers/2015/03/isis-twitter-census-berger-morgan/isis\\_twitter\\_census\\_berger\\_morgan.pdf](http://www.brookings.edu/~media/research/files/papers/2015/03/isis-twitter-census-berger-morgan/isis_twitter_census_berger_morgan.pdf).
4. ISIS, Terrorism Seen as Graver Threats Than Russia, Ukraine. Gallup. February 8, 2015. <http://www.gallup.com/poll/181553/isis-terrorism-seen-graver-threats-russia-ukraine.aspx>
5. Johnson A. Deviant and Pathological: What Do ISIS Extremists Really Want? NBC News. September 3, 2014. <http://www.nbcnews.com/storyline/isis-terror/deviant-pathological-what-do-isis-extremists-really-want-n194136>
6. Castillo M and Shoichet C. Was Arrested Teen on His Way to Join ISIS? CNN News. October 7, 2014. <http://www.cnn.com/2014/10/06/us/teen-isis-support-arrest/>
7. Why the US Government Is Trolling Terrorists on Social Media. CNN. April 3, 2014. <http://www.cnn.com/2014/04/18/world/jihadist-twitter-state-department-trolls-terrorists/>

8. Gladstone R. Twitter Says It Suspended 10,000 ISIS-Linked Accounts in One Day. The New York Times. April 9, 2015.  
<http://www.nytimes.com/2015/04/10/world/middleeast/twitter-says-it-suspended-10000-isis-linked-accounts-in-one-day.html>
9. Fighting ISIS on Twitter and the Internet. Lucky Troll Club. Web. April-May 2015.  
<http://luckytroll.club/daesh/ISISBlock.htm>
10. Reuters World News Headlines. September 23, 2014.  
<http://www.reuters.com/news/archive/worldNews?date=09232014>
11. Wang C. Scraping Tweets Using Twitter Stream API.  
<http://chengjun.github.io/en/2013/08/scraping-tweets-from-twitter/>

---

## Authors

Batsheva Moriarty is a senior at Harvard College concentrating in Computer Science with a secondary in Psychology. She has interned at Goldman Sachs and at BioCatch, a cyber-security start-up. In her free time she enjoys reading and traveling.

**Editor:** Latanya Sweeney

---

## Citation

Moriarty B. Defeating ISIS on Twitter. *Technology Science*. 2015092904. September 29, 2015. <http://techscience.org/a/2015092904>

---

## Data

Under review for data sharing classification. Data release available October 19.